HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

THE UNIVERSITY of York

WP 24/20

Health Inequalities and the Progressivity of Old-Age Social Insurance Programs

Jeroen van der Vaart; Max Groneck and Raun van Ooijen

August 2024

http://www.york.ac.uk/economics/postgrad/herc/hedg/wps/

Health Inequalities and the Progressivity of Old-Age Social Insurance Programs

Jeroen van der Vaart Max Groneck Raun van Ooijen*

May 31, 2024

Abstract

A well-established negative correlation exists between lifetime income and health and mortality risk. We quantify the welfare implications of living longer and using less LTC by higher incomes, implying higher lifetime retirement income and lower lifetime LTC cost. To this end, we model singles' and couples' consumption and saving behavior throughout the life cycle. Households face uncertain labor income at working age and uncertain and heterogeneous health and mortality across socioeconomic groups, so precautionary savings will differ across these groups. In addition, we assume that households value living and giving bequests to their heirs, implying a potential saving motive for bequests. We estimate the parameters of the model using unique administrative data from the Netherlands. Old-age insurance programs for retirement and LTC provision result in a substantial redistribution of welfare due to socioeconomic inequalities in LTC needs and mortality. The welfare effect amounts to 23.4% additional consumption after age 65 for the income-rich compared to those in the bottom lifetime income quartile. A large part of 22.2pp of the welfare gain for the richer households is explained by their strong preferences for leaving bequests: they have lower co-payments for LTC and more retirement income, which they spend on leaving a larger bequest upon death.

Keywords: Socioeconomic Inequalities, Long-term care and Mortality risk, Retirement Programs, Couples' Life-cycle Model.
JEL classification: D15, H55, I14, J14, J17

*Jeroen van der Vaart (A.J.van.der.Vaart@rug.nl), University of Groningen and Netspar; Rob Alessie, University of Groningen, and Netspar; Max Groneck: University of Groningen, AJSPH, and Netspar; Raun van Ooijen: University of Groningen, University Medical Center Groningen, AJSPH, and Netspar.
The authors thank Rob Alessie, Lieke Beekers, Hugh Gravelle, Jawa Izza, Minjoon Lee, Raymond Montizaan, Patrick Moran, Søren Leth-Petersen, and seminar participants at various seminars and conferences for valuable contributions.

1 Introduction

Health is strongly associated with socioeconomic status (Deaton, 2002; Chetty et al., 2016). This is a fundamental aspect of inequality in society with important implications for the progressivity of public health and social security policies for retired individuals (Poterba, 2014; Auerbach et al., 2017). As the income-rich live longer than the income-poor, they receive more years of social security benefits.[1] In contrast, better health may induce lower long-term care (LTC) needs for the income-rich, implying fewer years of costly out-of-pocket LTC expenditures, such as co-payments for nursing home use.[2] Thus, health inequalities could imply an unintended income-regressive redistribution.

This raises two important questions: What is the size of the welfare gain for households with higher socioeconomic status due to expecting to live longer and use LTC for a shorter time, and what mechanisms generate the gains? Such analysis should go beyond a conventional comparison of lifetime benefits and contributions and taxes (see, e.g., Goda et al., 2011a; Bosworth et al., 2016) because welfare consists of many other non-monetary factors, including the utility of consumption, bequeathing, and living longer (Bernheim, 1987). Yet, the size of the welfare gains and their mechanisms are not precisely pinpointed because a structural modeling approach allowing for counterfactuals is rarely applied.

This paper quantifies differences in the distribution of welfare due to socioeconomic inequalities in health. Furthermore, we investigate LTC co-payments and leaving bequests as mechanisms behind the differences. Bequests are particularly interesting because earlier work finds that wealthier households value these, and households can enlarge them when lifetime social insurance benefits are higher (De Nardi et al., 2010). We develop a life-cycle model of singles and couples where households value consumption, bequeathing, and living longer and are exposed to uncertain income during working age, and uncertain LTC use and mortality after that. Importantly, LTC use and survival risks differ exogenously by gender, marital status, and lifetime income quintiles to replicate the availability of informal care and the presence of socioeconomic differences in health.

[1] For socioeconomic inequality in mortality, see, e.g., Deaton (2002); Smith (2007); Chetty et al. (2016).

[2] For socioeconomic inequality in LTC use, see, e.g., Goda et al. (2011b); Jones et al. (2018); Rodrigues et al. (2018); Tenand et al. (2020a).

Our study focuses on the Netherlands with a generous and comprehensive LTC system (Bakx et al., 2023), including means-tested co-payments for nursing home care. We estimate the model using unique administrative data on income, assets, LTC needs, and mortality from 2006 to 2014. Finally, we shut down socioeconomic inequalities in health: we use the estimated model to compute how much consumption compensation each lifetime income quintile would require to be indifferent to being exposed to the LTC use and mortality risk of the bottom lifetime income quintile (cf. De Nardi et al., 2023). To see the impact of bequests and LTC co-payments, we remove them from the baseline model and re-compute the so-called consumption compensation equivalent.

Inspiration for our counterfactual and welfare measure comes from De Nardi et al. (2023). They use a structural life-cycle framework to quantify the lifetime cost of self-reported 'bad' health status for different initial health types. By assuming away the bad health shocks, they quantify the welfare cost of poor health for distinct health types. While we closely follow their approach, we conceptually differ as we shut down heterogeneity in health shocks rather than the health shock (uncertainty) itself.

The estimation proceeds in two steps. First, we estimate income, LTC, and mortality risk processes and calibrate the risk aversion parameter. Second, we include these health and income risks in a structural life-cycle model and estimate its key behavioral parameters: the subjective discount factor, consumption equivalence scale, the strength of the bequest motive, and the extent to which bequests are a luxury good. We estimate the parameters by matching simulated asset profiles to key aspects of the data, including asset holdings by marital status and lifetime income quintile. Also, we calibrate a parameter involving the Value of a Statistical Life, ensuring that households prefer living over death in utility terms (cf. Hall and Jones, 2007). After that, we use the estimated model to make counterfactual predictions.

Aligned with U.S. work, our findings identify leaving bequests as an important channel for the income-rich to save: we find the marginal propensity to bequeath to be unit value for every euro above a consumption level of 40 thousand euros. This saving motive almost exclusively involves households in the top lifetime income quintile; hence, bequests are luxury goods. Turning to differences between singles and couples, the estimated equivalence scale of consumption is 1.145 and lower than usually documented in the literature (see e.g., De Nardi et al., 2021), implying Dutch households can save more due to stronger economies of scale. Lastly, the estimated subjective discount factor of 0.960 reveals a moderate preference for current consumption.

In a subsequent counterfactual analysis, we remove socioeconomic differences in health risks. We find that moving from the counterfactual (no health differences) to the baseline (health differences exist) would increase consumption by 23.4% for the top income quintile after age 65. To stress the importance of our utilitarian framework, we report a monetary (accounting) gain of only 11.2%, driven mainly by more retirement benefits. Next, we assume away a preference for bequest saving and find that the welfare gain of 23.4% shrinks to 1.2% for the top lifetime income quintile. Hence, much of the welfare gain due to health inequalities stems from leaving larger bequests. Stated otherwise, increased bequest taxes could be a way to alleviate welfare gains due to living longer and using less LTC. If we remove (abolish) co-payments, the welfare gain remains 21.8%, implying that valuable bequests rather than co-payments explain the welfare gain.

Our paper builds on several different strands of literature. A recently developed literature quantifies the lifetime cost of (self-reported) bad health (see De Nardi et al., 2023, and the references therein). We merge this to the large macro-oriented literature that applies accounting and structural approaches to characterize the redistribution of old-age social insurance in a heterogenous-agent economy, where programs include Medicaid (see, e.g., De Nardi et al., 2016; Braun et al., 2017), Medicare (see, e.g., McClellan and Skinner, 2006; Bhattacharya and Lakdawalla, 2006), Social Security (see, e.g., Goda et al., 2011a; Fehr et al., 2013; Groneck and Wallenius, 2021), and co-payments for LTC (see, e.g., Wouterse et al., 2021). Auerbach et al. (2017) advocates a more holistic accounting approach that includes all old-age social insurance programs to report progressivity. Closest to our study, Bagchi (2019) and Jones and Li (2023) uses a structural life-cycle model to study the interaction between heterogeneous mortality rates and social security benefit formula reforms. We innovate this literature by going beyond self-reported health, mortality, and social programs; we also examine the contribution of heterogenous LTC use and preferences (bequests) to the welfare redistribution.

The current paper also contributes to the quantitative-micro literature on retirees' saving behavior. The desire to leave a bequest has received considerable attention as a potential explanation for why more affluent households retain high levels of wealth at very old ages, as in De Nardi et al. (2010), Lockwood (2018), Ameriks et al. (2020), and Nakajima and Telyukova (2023). However, the relative importance of this bequest saving and precautionary saving varies depending on the estimation strategy and data sample. De Nardi et al. (2010) finds an insignificant bequest saving motive, arguably because savings in the U.S. are fungible between high out-of-pocket medical expenditures and a bequest purpose (Dynan et al., 2004). Furthermore, the income-rich are under-represented in many surveys, including their Health Dynamics of the Oldest Old (AHEAD) data set. Lockwood (2018) instead finds a significant bequest saving motive by simultaneously fitting data on wealth and LTC insurance ownership, where the LTC insurance ownership acts as an exclusion restriction. The novelty of our paper is the use of data from a country where the need for precautionary saving against out-of-pocket medical expenditures is low and where the income-rich are well-represented in the administrative data.

Besides, we link to the scarce literature that empirically studies different saving behaviour by couples and singles within a life-cycle model (e.g., De Nardi et al., 2021). Like De Nardi et al. (2021), we view marital status shocks and their impact on LTC use as exogenous, indirectly capturing the relationship between informal care and LTC cost. It should be noted that for parsimony, we do not model the determinants of informal care; papers addressing such endogeneity stemming from altruistic and strategic informal care provision include Barczyk and Kredler (2018) and Ko (2022).

The paper is organized as follows. Section 2 presents the socioeconomic differences in health. Section 3 describes the life-cycle model. Section 4 provides the data and estimation procedure. Section 5 discusses the second-step estimation results. Section 6 performs the counterfactual health experiment. Section 7 discusses and concludes.

2 Socioeconomic Differences in LTC and Mortality

Before analyzing the welfare gain due to higher socioeconomic groups living longer and using less long-term care (LTC), it is crucial to examine how large these differences are. To quantify them, we follow van der Vaart et al. (2023) and use the same longitudinal data, socioeconomic status measure, and method as in to compute households' life histories on marital status, LTC use and death. In the analysis, we focus on 65+ individuals who are or were married at age 65; LTC use consists of institutional care use.[3] We observe 2,548,664 individuals and 1,487,109 households. See Appendix B.1 for a detailed description of the data and a summary of the estimation method.

Table 1: Life Expectancy and Long-term Care Use by Lifetime Income Quintiles

(a) Men	All	Bottom	Second	Third	Fourth	Top	ΔTop - Bottom
LE at age 65 (years)	18.9 (18.8;19.0)	16.4 (16.1;16.7)	17.6 (17.4;17.8)	18.4 (18.2;18.5)	19.2 (19.1;19.4)	20.0 (19.8;20.1)	3.6 (3.2;3.9)
LTC (years)*	1.9 (1.8;1.9)	1.8 (1.7;1.9)	2.0 (1.9;2.0)	2.0 (1.9;2.1)	1.9 (1.8;2.0)	1.7 (1.7;1.8)	−0.1 (−0.2;0.0)
Ever use LTC (%)	45 (44;45)	42 (40;43)	45 (44;46)	46 (45;47)	46 (45;47)	43 (42;44)	1 (0;3)
(b) Women							
LE at age 65 (years)	22.4 (22.3;22.4)	21.8 (21.5;22.0)	22.2 (22.0;22.4)	22.2 (22.1;22.4)	22.5 (22.3;22.6)	22.5 (22.4;22.7)	0.7 (0.4;1.1)
LTC (years)*	2.9 (2.9;2.9)	3.3 (3.1;3.4)	3.3 (3.2;3.4)	3.1 (3.0;3.1)	2.9 (2.8;2.9)	2.5 (2.5;2.6)	−0.7 (−0.9;−0.6)
Ever use LTC (%)	63 (62;63)	66 (65;68)	68 (67;69)	66 (65;67)	64 (63;64)	58 (57;58)	−9 (−10;−7)

Notes: Data: Dutch administrative records on individual and household marital status, gender, income, assets, institutional care use, and death between 2006 and 2014. The history of marital status dates back to 1995. Complete life histories on LTC use, deaths, and marital status are simulated according to van der Vaart et al. (2023). The presented numbers are population-averaged measures for the life cycle simulation of 100,000 households. We present the median estimates across 1,000 bootstrapped samples and the 2.5th and 97.5th percentiles between brackets. Appendix B.2 provides the goodness-of-fit between the simulated and empirical survival probabilities and LTC use rates by age, lifetime income, and gender.

Table 1 summarizes the remaining life expectancy (LE) and LTC use for men and women at age 65. We find opposite socioeconomic differences in LTC use and remaining life expectancy. Individuals within the top income quintile make less use of LTC but live longer. Men (women) within the top income quintile live 3.6 (0.7) years longer than their bottom income counterparts. Men (women) in the bottom income quintile use LTC for 1.8 (3.3) years on average, while this is 0.1 (0.7) years less for their top income

[3] Home-based care use is not a separate state because its co-payments and, thus, redistributive effects are very limited in the Netherlands (Tenand et al., 2020b).

counterparts. Although the difference in LTC use is small for men, it is substantial for women and amounts to 26% of their average duration of using LTC. The larger difference for women can partially be explained by the fact that they often outlive their partner and thus lose a potential source of informal care.

3 Life-Cycle Model

We develop a life-cycle model with uncertain LTC use and mortality to quantify the welfare gain of the higher lifetime income quintiles using less LTC and living longer. At every age $t \in \{25, 26, ...100\}$, a household maximizes lifetime utility by choosing total consumption expenditures c and savings a. The savings also determine the bequest that is left upon the death of the last household member. Households derive utility from consumption, leaving bequests, and being alive (independent of consumption). For tractability, we assume that household members are the same age such that a single age suffices to characterize the household.

A household has one of the following family statuses (f): a couple, single woman, or single man. Households enter the model as a couple initially and remain a couple until retirement at age 65, so there is no divorce or widowhood. Also, we assume no use of LTC before age 65 because of low likelihood.[4] After age 65, survival and use of LTC become uncertain, and couple households can become a single woman or single man household.

3.1 Preferences

The per-period CRRA utility functions of couples (C) and singles (S) are given by:

$$u^C(c) = 2 \cdot \frac{\left(\frac{c}{\eta}\right)^{1-\sigma}}{1-\sigma} + \bar{b}, \quad \text{and} \quad u^S(c) = \frac{c^{1-\sigma}}{1-\sigma} + \bar{b}, \quad \sigma \geq 0, \quad 1 \leq \eta \leq 2, \quad \bar{b} \geq 0,$$

where the parameter $\sigma \geq 0$ reflects the level of risk aversion.

Following the literature (De Nardi et al., 2021), we allow couples to benefit from economies of scale. Partners can pool their income and consume many goods jointly. η determines the extent to which households benefit from economies of scale. $\eta < 2$ features

[4] At age 65, only 1% of the sample uses LTC.

economies of scale: each couple member consumes $\frac{c}{\eta}$ units while this would be $\frac{c}{2} < \frac{c}{\eta}$ if the two do not share a household (Browning et al., 2013).

Following De Nardi et al. (2023), we introduce scaling parameter $\bar{b} > 0$. This parameter is crucial when examining the welfare implications of altered life expectancies because households could attach value to the 'invisible' good of being alive that goes beyond consuming and bequeathing, e.g., the happiness of being alive.[5] In our model, risk-averse households would reach higher utility when being dead because $u^S < 0$ and $u^C < 0$ and utility from death is zero. We assume that utility from being alive is higher, thus calibrating a $\bar{b}$ yielding non-negative utility in any state when alive: $u^S \geq 0$ and $u^C \geq 0$.

The household derives utility $\mathcal{B}(a)$ from leaving bequest a. Following De Nardi (2004):

$$\mathcal{B}(a) = \frac{\phi}{1-\phi}^{\sigma} \cdot \frac{\left(\frac{\phi}{1-\phi} \cdot c_a + a\right)^{1-\sigma}}{1-\sigma} \quad \text{if } \phi \in (0,1),$$

$\mathcal{B}(a) = c_a^{-\sigma} \cdot a$ if $\phi = 1$ and $\mathcal{B}(a) = 0$ if $\phi = 0$, which De Nardi (2004) introduced to be consistent with wealth concentration among the wealthiest households in the U.S.. c_a is the consumption level below which households, under perfect certainty, will not leave a bequest (Lockwood, 2018). $c_a > 0$ implies bequests to be luxury goods. If households' wealth meets threshold c_a, ϕ is the share of excess wealth spent on a bequest: higher ϕ increases marginal utility from bequeathing relative to marginal utility from consuming.

3.2 Sources of Uncertainty

An important empirical artifact to be replicated is heterogeneity in asset holdings. A source for heterogeneity is uncertainty, forcing households to make precautionary savings (Carroll, 1997). We have uncertain health, family status, and income in our model.

Use of LTC and survival After age 65, exogenous health and family status shocks occur. The health of the husband and wife, h^m and h^f, evolve jointly and can differ between them ($h^m \neq h^f$). h^m and h^f take three values: a household member does not use public institutional care ($i = 1$), uses public institutional care ($i = 2$), or is dead ($i = 3$).

[5]In the literature, this parameter is used to compute the Value of a Statistical Life, i.e., the price that a population is willing to pay to prevent one certain death in the current period (see, e.g., Hall and Jones, 2007; St-Amour, 2022). This statistic is outside the scope of our study.

LTC use induces co-payments (out-of-pocket expenditures) $m(y, a, h^m, h^f)$ that depend on income (y) and assets (a); these are paid to the government. We assume that LTC needs are homogenous across institutionalized individuals, so co-payments do not depend on the severity of the need for care.

We assume a Markovian process, so transition probabilities depend on the health and survival statuses of the preceding period: h_t^m and h_t^f. Survival status of a spouse controls for the potential availability of informal care. Furthermore, the health transition probability depends on lifetime income I and age t. Health transition probability π is:

$$\pi_{k,l}^{i,j}(t, I) = \mathbb{P}(h_{t+1}^m = k, h_{t+1}^f = l \mid h_t^m = i, h_t^f = j, t, I) \quad \text{with: } (i, j, k, l) \in \{1, 2, 3\}.$$

The death probability of the households is as follows:

$$\pi_{3,3}^{i,j}(t, I) = \mathbb{P}(h_{t+1}^m = 3, h_{t+1}^f = 3 \mid h_t^m = i, h_t^f = j, t, I) \quad \text{with: } (i, j) \in \{1, 2, 3\}.$$

Life-cycle income: age 25 to 65 Exogenous income shocks happen during working life, reflecting the presence of labor supply shocks and health shocks. To save on the state space, we assume that these income shocks occur at the household level. Following the standard literature (Storesletten et al., 2004; French, 2005), household income dynamics follow an AR(1) process:

$$\begin{aligned} y_t &= \min(\widetilde{y}_t; \underline{y}) \qquad (1) \\ \widetilde{y}_t &= \alpha_t \cdot \exp(\theta) \cdot \exp(\eta_t) \cdot \exp(\epsilon_t) \\ \eta_t &= \rho \cdot \eta_{t-1} + u_t \\ \theta &\sim \mathcal{N}(0, \sigma_\theta^2); \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2); \quad u_t \sim \mathcal{N}(0, \sigma_u^2); \quad \eta_{24} = 0, \end{aligned}$$

where y_t is pre-tax household income, including income from labor, capital, and social insurance. α_t a deterministic age effect. θ is a fixed (labor) productivity effect. η_t is a persistent shock. ϵ_t is a transitory shock, in part reflecting transitory health shocks. η_{24} is the initial level of the persistent income part. $\underline{y}$ is a government-provided income floor.

Life-cycle income: age 65 and older Households receive retirement income $y_t = SS(f) + DB_{65}(f)$ consisting of a part independent of the income history, SS, and a part DB_{65}, whose defined benefit formula depends on the income history $\{y_s\}_{s=25}^{64}$. Income depends on the family structure because it becomes a smaller survivor benefit upon widowhood. Retirement income is stochastic due to random shocks in the income history until age 65 and time-varying family status.

3.3 The Government

The government provides income and LTC insurance after retirement by providing a first pillar pension and (partially) covering institutional care costs. Households pay mandatory for this insurance via dedicated taxes $\tau_{SS}(y)$ and $\tau_L(y, f, t)$. Moreover, co-payments $m(y, a, h^m, h^f)$ finance LTC use. Lastly, households pay a general income tax $\tau_G(y, f, t)$. We specify the functional forms of tax function τ in Appendix B.6. We specify m in Section 4.

Government revenues and costs in the model do not necessarily balance, which we ensure with additional lump-sum transfers Tr_{SS} and Tr_{LTC}. Appendix A.1 describes the procedure for how the government sets these transfer levels.

3.4 Optimization Problem

The timing is as follows: at the beginning of the period, households observe their state variables $\aleph$ that are relevant to their decision-making. The household obtains interest rate r on assets a, obtains income y, pays taxes τ and co-payments m, and makes the government-balancing transfers Tr. Then, households 'optimally' consume or save the remaining assets based on state vector $\aleph$. Lastly, a survival and LTC use shock hits. If the final household member has died, any remaining assets go to the household's heirs (we assume households value their gross bequest and, therefore, ignore bequest taxes).

The state vector, $\aleph$, represents variables that are commonly observed by the household at the beginning of each period t:

$$\aleph_t^W = (a_t, \theta, \eta_t, \epsilon_t, \mathrm{DB}_t, t)' \qquad \text{(if } t < 65\text{)}$$

$$\aleph_t^R = (a_t, \mathrm{DB}_{65}, f_t, h_t^m, h_t^f, t)', \qquad \text{(if } t \geq 65\text{)}$$

where after age 65, retirement income replaces stochastic income, and family status f_t and health statuses h_t^m and h_t^f become uncertain. DB_t is the pension accrual until age t.

Note that all variables are known before deciding consumption c_t and next period's assets a_{t+1}, so we can recursively write the household's problem. Denote β the subjective discount factor. The household's value function at age t is:

$$V_t(\aleph_t^W) = \max_{c_t, a_{t+1}} \; u^C(c_t) + \beta \cdot \mathbb{E}[V_{t+1}(\aleph_{t+1}^W) \mid \aleph_t{}^W]. \qquad \text{(if } t < 65)$$

$$\begin{aligned} V_t(\aleph_t^R) = \max_{c_t, a_{t+1}} \; & u^f(c_t) + \beta \cdot \left(1 - \pi_{3,3}^{i,j}(t, I)\right) \cdot \mathbb{E}[V(\aleph_{t+1}^R) \mid \aleph_t{}^R] \\ & + \beta \cdot \pi_{3,3}^{i,j}(t, I) \cdot \mathcal{B}(a_{t+1}), \qquad \text{(if } t \geq 65) \end{aligned}$$

subject to a budget constraint and no-borrowing constraint, defining next period's assets:

$$a_{t+1} = (1 + r) \cdot a_t + y_t - \tau_G - \tau_{SS} - \tau_L - m_t - \mathrm{Tr}_{SS} - \mathrm{Tr}_{LTC} - c_t \geq 0.$$

The dynamic optimization problem after age 65 is different due to health uncertainty. A household survives into the next period with probability $1 - \pi_{3,3}^{i,j}(\cdot)$, and then faces the optimization problem again (V_{t+1}). With probability $\pi_{3,3}^{i,j}(\cdot)$, the household leaves a bequest with utility flow $\mathcal{B}(a_{t+1})$. Also co-payments for LTC use might occur ($m_t \neq 0$). We discuss the numerical implementation in Appendix A.2 to A.4.

As will be later important for our counterfactual analyses, health h_t^f and h_t^m impact the decision problem both via the utility function and budget constraint. The survival probabilities are lower when using LTC, implying that future consumption is more heavily discounted and households save less for future consumption. Health ambiguously affects the decision problem via the co-payments. On the one hand, co-payments for LTC limit the available budget for consumption, inducing the household to precautionary save. On the other hand, a co-payment depends on assets and puts a penalty on saving.

4 Data and Estimation Procedure

For estimation, we use administrative data from Statistics Netherlands that is available under restricted access. We can merge different data sets within the secured environment based on a unique individual and household identifier. Data come from multiple sources and registries: tax files (income and assets), municipal population registries (marital status, gender, birth year, and age), and a registry on institutional care use and deaths.

We use a two-step strategy similar to Gourinchas and Parker (2002) and De Nardi et al. (2010) to estimate the unknown parameters of our life-cycle model. In the first step, we estimate parameters directly from the data and denote them by $\boldsymbol{\chi}$. For example, we estimate parameters for the health and income processes from the administrative data. Without estimation, we tailor the pension and LTC use system to the Dutch setup 2006-2014. In addition, we fix the risk aversion and interest rate to $\sigma = 3$ and $r = 2\%$, values commonly used and found in life-cycle studies (see, e.g., De Nardi et al., 2010).

Given the parameters and shock processes from the first stage, we apply the method of simulated moments to estimate the remaining parameters. We minimize the sum of squared differences between empirical and simulated moments of the asset distribution. The parameters to estimate are the subjective discount factor, bequest utility parameters, equivalence scale of consumption, and government-balancing transfers: $\boldsymbol{\delta} = (\beta, \phi, c_a, \eta, \text{Tr}_{SS}, \text{Tr}_{LTC})'$. After estimating all the parameters, we calibrate $\bar{b}$, i.e., the scaling parameter for the utility of surviving households.

4.1 First-Step Calibration and Estimation

Use of LTC and survival We estimate the health transition matrix using our simulated sample on household use of LTC from Section 2. We convert the life histories from continuous time to discrete time (an age period of one year), and compute transition probabilities accordingly. LTC use is assumed to be used throughout the entire age period and yearly costs the government €58,500 per user. The model is estimated using daily reported deaths, institutional care use, and marital status between 2006 and 2014. See Appendix B.1 for a detailed description of the data and a summary of the estimation method, including the computation of lifetime income (quintiles).

Co-payments for LTC use In the Netherlands, households make a co-payment to finance the use of LTC. The co-payment depends on the asset level a, household income y, and health statuses h^m and h^f. Households pay a low-rate or high-rate co-payment depending on the LTC used by the household members (h^m and h^f). The low-rate co-payment applies to couples with only one LTC user. The high-rate co-payment applies to singles and households with two LTC users:[6]

$$m(y,a,\cdot)\begin{cases} \max[1,900,\ \min[9,800,\ \ 0.125\cdot(y+0.04\cdot a)]] & \text{(low co-pay)} \\ \max[0\qquad,\ \min[27,000,\ \ 0.75\cdot(y_{AT}+0.04\cdot a-4,500)^{+}]] & \text{(high co-pay)} \end{cases}$$

The main difference between the two co-payment types stems from the cap on co-payments, €9,800 vs. €27,000, and the co-pay rate on income: 0.125 vs. 0.75. Also, note that contrary to low-rate co-payments, high-rate co-payments depend on income after taxes $y_{AT} = y - \tau_G - \tau_{SS} - \tau_L$. Lastly, 4% of the assets contribute to co-payments, implying endogenous co-payments in the model. In 2013, a policy change imposed an additional 8% of the assets to count for the co-payments. However, we stick to the 4% because that spans most of our sampling window (2006-2014).

Life-cycle income: age 25 to 65 To estimate the income shock process (y_t), we use granular income data available for a representative sample of about 1% of the households (the IPO sample). In this sample, we have information on the distinct categories that comprise household income (the IPO sample), including taxes and private and public pension benefits. The data are available for a longer period than the data for the health processes: 2001-2014. A longer sampling window is indispensable when estimating the persistence, i.e. longstanding effects, of income shocks.

We observe pre-tax income aggregated to the household level, including social transfers and pension income. This income definition also includes taxes for first pillar pension income and LTC provision, and a general income tax but excludes other dedicated taxes, e.g., for unemployment insurance. We only include the income of the household head and the partner (if applicable) and exclude the income of other household members. The variables are normalized to base year 2015 with the Consumer Price Index.

To abstract from early retirement decisions and schemes, we restrict our sample to couples whose oldest member is born after 1949 and whose primary income source is not retirement income. Further, we only include income above the government-provided

[6]We keep the formula simple for computational reasons, but the system is more complex in practice. Income and assets are measured with a two-year lag, implying we would have two additional state variables in our model. A low co-pay rate applies for the first four months of an institutional stay, which we cannot measure with our model specified at the year level instead. Also, there is an asset exemption of about €21,000 and €42,000 for singles and couples, but we follow Wouterse et al. (2021) and (partially) replace this with a general exemption of 0.75× €4,500 for the high co-payment.

safety net (welfare level): $y > \underline{y} = €15,600$ (2010-level). $\underline{y}$ is a government-provided income floor, equivalent to a consumption floor, as in, e.g., De Nardi et al. (2023).

We follow Storesletten et al. (2004) for the estimation of the income shock process. We estimate the age effect α_t and productivity effect θ_i by running a fixed effects regression of log income on age dummies (one for each $\log(\alpha_t)$) and a household fixed effect (θ_i):

$$\log(y_{it}) = \log(\alpha_t) + \theta_i + \eta_{it} + \epsilon_{it}, \tag{2a}$$

where i indexes a household and t the age of the oldest household member.

Ideally, our household-specific estimate $\widehat{\theta}_i$ excludes birth year effects. To wash out present cohort effects, we run the following OLS regression of the predicted productivity effects on birth year dummies (cf. French, 2005; De Nardi et al., 2023):

$$\widehat{\theta}_i = \overline{c} + \overline{\theta}_c + \widetilde{\theta}_i, \;\; c \in \{1951, .., 1990\}, \tag{2b}$$

where $\overline{c}$ is the cohort effect of birth year 1950, $\overline{c} + \overline{\theta}_c$ is the cohort effect for birth years 1951-1990, and residual $\widetilde{\theta}_i$ is the household productivity effect excluding a cohort effect. We use $\widetilde{\theta}_i$ as the household-specific productivity effect.[7]

Next, we estimate the parameters of the income shock $\theta_i + \eta_{it} + \epsilon_{it}$: ρ, σ_θ, σ_u, and σ_ϵ. To this end, we construct the empirical auto-covariance matrix of the predicted residuals of $\tilde{\theta}_i + \eta_{it} + \epsilon_{it}$ from (2a) and (2b), and match them to the auto-covariances implied by equation (1). Appendix B.4 further explains the GMM procedure and shows the fit.

Table 2: Parameters of the AR(1) Income Process

Parameter:	ρ	σ_θ	σ_u	σ_ϵ
	0.966	0.184	0.131	0.166
	(0.004)	(0.028)	(0.008)	(0.003)

Estimates for married households whose oldest member is younger than age 65 and born after 1949. Data from IPO 2001-2014: 77,118 households and 534,006 panel-year observations. Standard errors in parentheses.

Table 2 provides the results on the income shock. The estimated parameters align with results in the literature (Storesletten et al., 2004; Karahan and Ozkan, 2013; Blundell et al., 2015; Paz-Pardo and Galves, 2023). This also holds for the high income persistence $\rho = 0.966$ we estimate: income shocks have longstanding effects.

[7] Appendix B.3 shows the model estimates for the age profile $\{\overline{c} + \log(\alpha_t)\}_{t=25}^{64}$.

Life-cycle income: after age 65 In the Netherlands, first pillar pension income is independent of income history $\{y_s\}_{s=25}^{64}$ but linked to minimum wage w. For couples, the benefit level is minimum wage ($SS = w$). For singles, the benefit level is 70% of the couple's benefit ($SS = 0.7w$). As minimum wage we take the 2010-value: $w =$ €$18,240$.

A household in the model is also entitled to a second pillar pension benefit DB_{65}, which is linked to the history of income shocks $\{y_s\}_{s=25}^{64}$. In practice, the first and second pillar aim to replace about 75% of the average individual-earned income or obtained disability insurance income (Knoef et al., 2017).[8] We assume the same replacement rate and benefit formula at the household level. The second pillar pension income is only accrued over the income y_t that exceeds $\frac{100}{75} \cdot SS$ because social security benefits are sufficient to replace the income below this level. The evolution of the second pillar pension benefit is:

$$\mathrm{DB}_{t+1} = \mathrm{DB}_t + \frac{1}{40} \cdot 0.75 \cdot \min\left(y_t - \frac{100}{75} \cdot SS \;\; ; \;\; 0\right) \quad \text{if } t \leq 64,$$

where the factor $\frac{1}{40}$ makes sure we take a 40-year average of pre-tax household income.

Together, first and second pillar pensions compose income after retirement ($t \geq 65$):

$$y_t(\mathrm{DB}_{65}, f) = y_t(\{y_s\}_{s=25}^{64}, f) = \begin{cases} w + \mathrm{DB}_{65}, & \text{if} \quad f = \text{couple} \\ 0.7w + \mathrm{rr}_w \cdot \mathrm{DB}_{65}, & \text{if} \quad f = \text{single woman} \\ 0.7w + \mathrm{rr}_m \cdot \mathrm{DB}_{65}, & \text{if} \quad f = \text{single man.} \end{cases}$$

If a spouse dies, rr_w and rr_m convert a couple's pension benefit into a widow(er)'s pension benefit. Using the IPO data, we find $\mathrm{rr}_m = 0.93$ (SE: 0.001) and $\mathrm{rr}_f = 0.55$ (SE: 0.005). In line with our earlier work van der Vaart et al. (2020), we report $\mathrm{rr}_m > \mathrm{rr}_f$ implied by that men were the prime earner in the households and pension benefits mostly accrued to them. Appendix B.5 contains the estimation details.

A crucial variable in our model is lifetime income quintile I, which determines the health risks after retirement. We take DB_{65} as the model-equivalent level of lifetime income, which is exogenous because households do not decide on labor supply. Consequently, we can compute the quintiles of the distribution of DB_{65} without running the life-cycle model. Next, when running the life-cycle model, we use the quintiles and realization of DB_{65} to assign households a lifetime income quintile group.

[8] We keep the formula simple for computational reasons, but the system is more complex in practice.

Taxation We estimate the tax function $\tau_{SS}(y)$, $\tau_L(y,f,t)$ and $\tau_G(y,f,t)$ by regressing observed tax amounts in the IPO on household income according to a log-linear and sigmoid specification. We apply non-linear least squares estimation and estimate the functions separately for households below and above age 65 and for single and married households. Appendix B.6 reports the specifications and estimates.

Remaining calibrations Table 3 displays the remaining first-stage parameters.

Table 3: Other First-Step Parameters

	Symbol:	**Value:**	**Source:**
Relative risk aversion	σ	3	Several empirical studies[1]
Interest rate	r	0.02	The average interest rate on savings 2006-2014[2]
First pillar pension benefit	w	€18,240	2010-level
Yearly LTC cost per user (€)	$\mathrm{LTC}_{\mathrm{cost}}$	€58,500	van Ooijen et al. (2015)

[1] See estimates by Cagetti (2003); De Nardi et al. (2010); Lockwood (2018)
[2] See DNB Statistics (2023): https://www.dnb.nl/statistieken/dashboards/rente/, [August 7, 2023]

4.2 Second-Step Estimation

In this step, we apply the method of simulated moments (MSM) estimation to match asset moments in the administrative data with moments simulated with the life-cycle model (see, e.g., De Nardi et al., 2010; Lockwood, 2018; De Nardi et al., 2021). Using our estimated first-stage parameter vector $\boldsymbol{\chi}$, we try to find preference vector $\boldsymbol{\delta} \in \boldsymbol{\Delta}$ that yields model-generated asset profiles that 'best match' observed asset profiles. The matching happens according to standard generalized method-of-moments (GMM) techniques.

For the empirical moments, we use the same population and lifetime income quintiles that we used to compute the health process in Section 2, i.e., households whose members are aged older than 65 and were married at age 65. Following seminal work on the elderly's asset holdings (De Nardi et al., 2010; Ameriks et al., 2020; Nakajima and Telyukova, 2023), we take net worth as our measure of wealth. This is the total assets reduced by mortgages and other debt. Total assets are defined as the sum of the values of checking and savings accounts, risky assets (stocks and bonds), business wealth, the owner-occupied house, other real estate, and other assets such as cash-on-hand. The value of risky assets is normalized with the Amsterdam Exchange close index (AEX) on 31/12/2014, the owner-occupied house and other real estate with the house price index (base 2015), and debt and amounts deposited in checking and savings accounts with the Consumer Price Index (base 2015).

To prevent an overly complex model, we do not separately treat financial wealth and net housing wealth, i.e., the total value of real estate minus outstanding mortgage debt. The co-payments are, however, based on financial wealth and exclude housing wealth in practice. We assume that households liquidate their housing wealth (sell their house) once they enter a public care institution. Hence, net worth and financial assets coincide.

We base our estimator on the age profile of the median net worth of married and single individuals between ages 65 and 100 by lifetime income quintile I, implying $2 \times 36 \times 5 = 360$ moment conditions. We do not consider matching means (cf. De Nardi et al., 2010) because these empirical moments are sensitive to outliers, thereby driving estimation results. Furthermore, we restrict the analysis to matching the asset distribution after age 65 because our studied welfare effects primarily occur after this age.

However, akin estimating the income processes before age 65, we must first deal with cohort effects to observed asset profiles. We similarly account for this as specifications (2a) and (2b) do for the income process. To stay as close as possible to the 1950 cohort for which we estimated the income process, we made the assets representative for a reference group of households born between 1945 and 1949. Appendix B.7 provides details about how we econometrically deal with the cohort problem of assets.[9]

We compute the moments also for our simulated sample and compare them with the data moments using the objective function:[10]

$$\sum_{k=1}^{K=360} \left[\left(M_k^d - M_k^s(\hat{\chi}, \boldsymbol{\delta}) \right)^2 \right],$$

[9]The regressions involve the logarithm of assets, so we only keep non-negative assets. Furthermore, the regression is prone to outliers, so we drop assets above €2,500,000. We drop 0.9% of the households and 2.6% of the panel-year observations because of these restrictions.

[10]Instead of matching medians directly, existing work (e.g. Cagetti, 2003) looks at how many households in the observed population have assets below the simulated median, which is ideally 50%. This means that at each iteration, we would have to use our administrative data to determine how many individuals have assets below the group-specific simulated median, which is computationally expensive. That condition and our condition are equivalent at the true value $\boldsymbol{\delta}$ so we choose our current approach.

with $K = 360$ moments, and where M_k^d and M_k^s are the k-th data and simulated moment.

Our estimator $\hat{\boldsymbol{\delta}}$ minimizes this quadratic distance between the empirical and simulated data moments. We do not weight each moment with an asymptotically optimal weight matrix, implying we have less efficient estimates. Instead, efficient estimates would follow from taking the densities evaluated at the median as weights (Powell, 1994), but estimating these weights is computationally too expensive.[11]

The procedure can be summarized as follows. We first estimate asset profiles from the administrative data. Second, we estimate the unknown parameters for the first stage. Then, we take the first-stage calibrations $\hat{\boldsymbol{\chi}}$ and a given parameter value $\widetilde{\boldsymbol{\delta}}$ and run the life-cycle model. We store the decision rules of the life-cycle model. We know the steady-state distribution of individuals over the state variables and can compute the simulated asset moments from that (see Appendix A.4 for the computation of the distribution). Hereafter, the value of the objective function is computed. Lastly, we compute a new 'optimal' preference vector using a Gauss-Newton regression and repeat the procedure until parameter vectors of two consecutive iterations are arbitrarily close. See Appendix C.1 for the computation of the standard errors.

Lastly, we calibrate $\overline{b}$, a crucial parameter when examining the welfare implications of shortening and extending lifespans (cf. Hall and Jones, 2007). Our additive speciation implies that $\overline{b}$ does not depend on the consumption and saving decision, so we do not have to jointly estimate this parameter with the other preference parameters, but rather calibrate it conditionally upon them. We tailor the parameter to the group that has the lowest-per-period utility in our population: retired singles without private pensions ($\text{DB}_{65} = 0$). We set $\overline{b} = -\frac{\underline{c}^{1-\sigma}}{1-\sigma} = 0.3114$, where $\underline{c} = 0.7w =$ €$12,768$ is their consumption level (in 0000s €) and implying this group has zero utility from consumption. In a similar spirit, De Nardi et al. (2023) used an estimated consumption floor to pinpoint $\overline{b}$. Because we tailor $\overline{b}$ to the lowest consumption level, our estimated welfare gains from living longer will be a lower bound to the true effect.

[11] We also tried inverse-variance weighting (cf. Altonji and Segal, 1996). However, this implied non-sensible estimates as there are extremely large weights for low compared to high lifetime income quintiles.

4.3 Model Identification

β is identified by the shape of the age profile on assets: higher β implies a stronger preference for future consumption and thus more saving. In addition, the Euler equation provides intuition for the identification of preference parameters on bequests ϕ and c_a, and equivalence scaling η in our model. To see how this works for η, suppose a simple model without bequests of a married household in period t, that will be not be married in period $t+1$ anymore. If the sole uncertainty is death, then the Euler equation implies the following consumption growth:

$$\log\left(\frac{c_{t+1}^S}{c_t^M}\right) = \log\left(c_{t+1}^S\right) - \log\left(c_t^M\right) = -\log(\eta) + \frac{1}{\sigma}\cdot\left(\log(\beta) + \log(1-\pi_{3,3}) + \log(R) - \log(2)\right),$$

where c_{t+1}^S is consumption when single, and c_t^M is consumption when married. Here, higher η (less economies of scaling) implies more consumption spending c_t^M when married, so lower savings when married. Hence, we identify η by comparing asset levels of married and single households of a given lifetime income quintile at two consecutive ages.

Also, the Euler equation shows a complication when having to estimate β and σ. Their joint effect on savings would be $\frac{1}{\sigma}\cdot\log(\beta)$, making it impossible to separately identify the two when studying a given asset level. Therefore, we follow Ameriks et al. (2011) and fix $\sigma = 3$, a value common in retirement-savings literature (e.g. De Nardi et al., 2010).

Lastly, to see how the bequest parameters are identified, we consider a single household that knows to die next period, does not subjectively discount utility from consumption c, and obtains utility from leaving a bequest a ($c_a > 0$ and $\phi \in (0,1)$). Assume that the household has cash-on-hand μ, then the decision problem is:

$$\max_{c,a}\ u^S(c) + \mathcal{B}(a) = \max_{c,a}\ \frac{c^{1-\sigma}}{1-\sigma} + \bar{b} + \frac{\phi}{1-\phi}^{\sigma}\cdot\frac{\left(\frac{\phi}{1-\phi}\cdot c_a + a\right)^{1-\sigma}}{1-\sigma}, \quad s.t.\quad \mu = a + c.$$

The Euler equation with bequests is:

$$c^{-\sigma} = \frac{\phi}{1-\phi}^{\sigma}\cdot\left(\frac{\phi}{1-\phi}\cdot c_a + a\right)^{-\sigma}, \quad s.t.\quad \mu = a + c \rightarrow$$

$$c = c_a + (1-\phi)\cdot\mu \text{ and } a = \phi\cdot(\mu - c_a),$$

so in optimum households equate the marginal utility from bequests and consumption.

Increasing c_a one-to-one increases consumption c, and one-to-one decreases the bequest size a. c_a is thus a terminal wealth level that must be met before a household intends to leave a bequest (bequests are luxury goods).[12] The likelihood of meeting this criterion is larger for higher lifetime income quintiles, from whose terminal assets we identify c_a. Furthermore, ϕ is the share of excess wealth they leave as a bequest. We identify ϕ by comparing the steepness of the asset profile for this group with $\mu > c_a$ compared to the groups with insufficient wealth $\mu \leq c_a$, i.e. groups with low lifetime income.

5 Second-Step Estimation Results

Figure 1 shows the empirical and simulated moments for our closest match. For exposition, we connect the moments with a line. Overall, we have a reasonable fit: we match the positive correlation between the level of assets and lifetime income quintile and the asset decumulation pattern after age 65. We also mimic the empirical artifact that households in the top income quintile die with substantial assets, i.e., leave a bequest. Our model is less capable of matching the low asset holdings for the bottom and second income quintile, which could be explained by that these groups contain relatively many hand-to-mouth consumers and have lower discount rates (Cherchye et al., 2023). However, introducing heterogenous preferences would make it less clear where a welfare redistribution stems from and abstracts from the standard in the retiree's saving literature that we stick to, i.e. a parsimonious model with homogenous preferences (De Nardi et al., 2010; Ameriks et al., 2020). Yet, the general picture of asset profiles seems to be reproduced by our MSM estimation, making us confident in using our estimated life-cycle model for further inference.

[12] In the model, the survival probability is below unit value, so c_a refers to annuitized cash-on-hand rather than the level of cash-on-hand.

Figure 1: Empirical and Simulated Assets by Lifetime Income Quintile and Marital Status

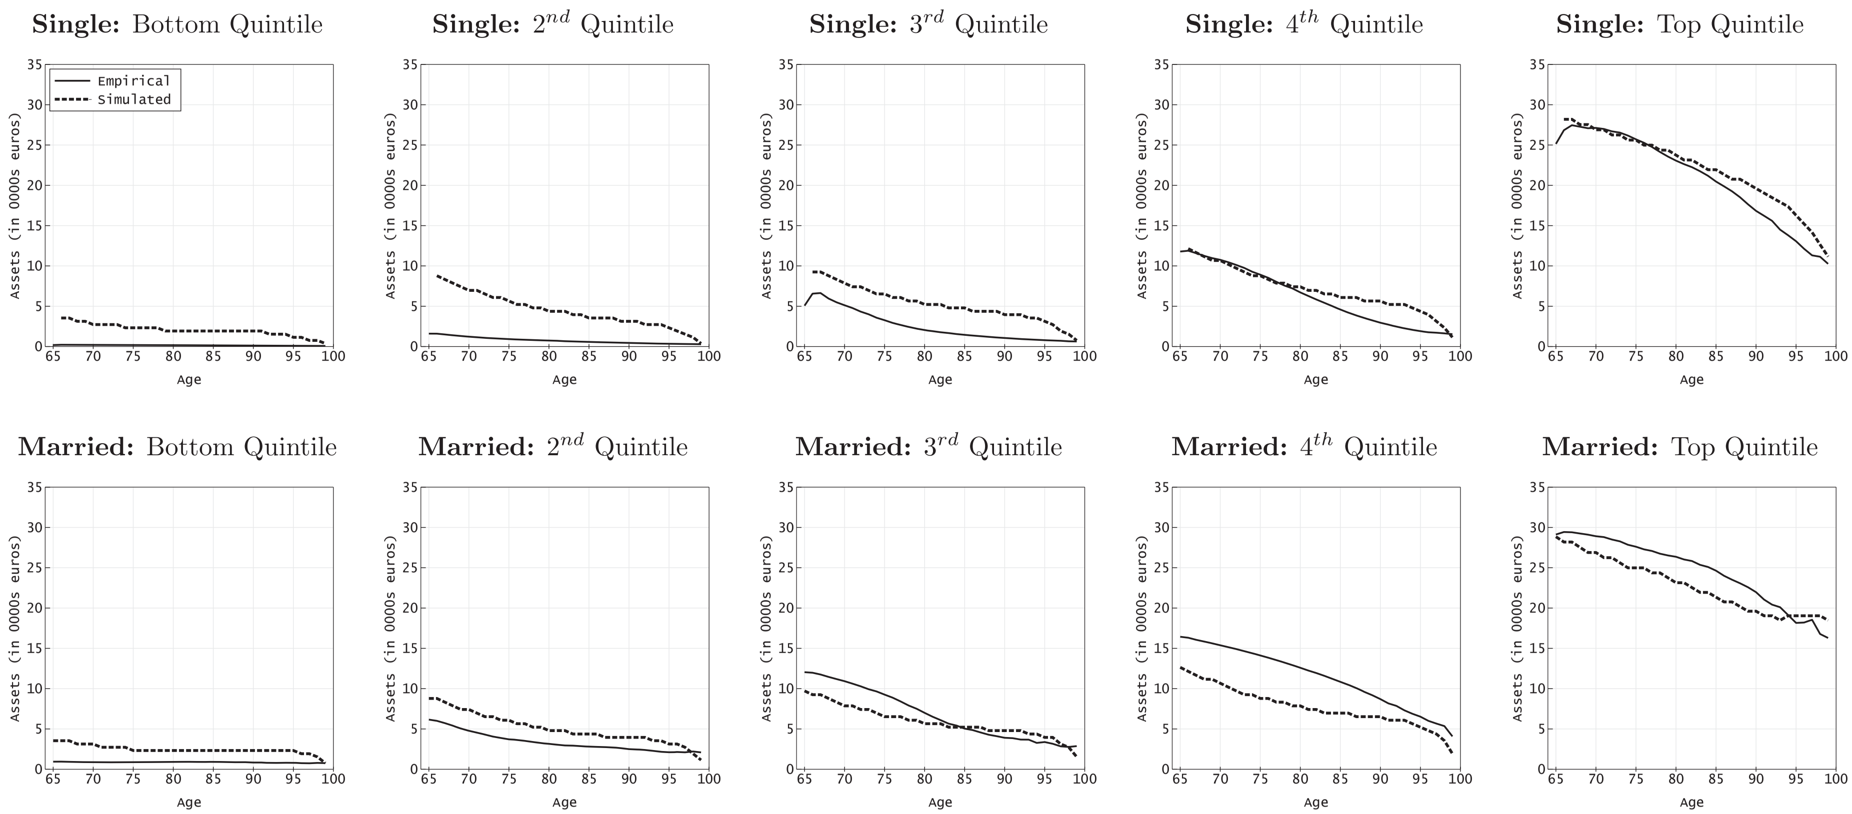

Table 4 presents the results of our preference parameter estimation. We estimate $\hat{\beta} = 0.960$, implying that households have a moderate preference for current over future consumption. The estimated bequest utility indicates a strong saving motive, where bequests are luxury goods ($\widehat{c_a} = 40,672 > 0$). We find the extreme case of $\widehat{\phi} = 1$, implying a linear bequest function, and all excess wealth is put into a bequest and not consumed. A high bequest propensity ($\phi > 0.88$) is common in the revealed preferences literature (De Nardi et al., 2010; Lockwood, 2018; De Nardi et al., 2023), while the stated preference literature finds lower values ($\widehat{\phi} > 0.48$, see, e.g. Ameriks et al., 2020). Our threshold consumption level $\widehat{c_a}= 40,672$ is close to De Nardi et al. (2010), who report $\widehat{c_a}= 34,000$ and slightly higher than other related studies (Lockwood, 2018; De Nardi et al., 2023).

Table 4: Estimated Structural Parameters

Discount factor	Bequest utility		Equivalence scale	Government transfers	
β	c_a	ϕ	η	Tr_{SS}	Tr_{LTC}
0.960 (0.00002)	40,452 (1.03844)	1.000 (0.00054)	1.145 (0.00010)	783.58	-433.05

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors in parentheses. The data contain 1,471,858 households and 11,471,725 panel-year observations.

We find an equivalence rate of $\widehat{\eta} = 1.146$, which is lower than the commonly applied and estimated OECD-modified equivalence scale of 1.5 (for the life-cycle model estimate, see, e.g., De Nardi et al., 2021). Lower equivalence scales are, however, also reported in the consumption-expenditure literature (see, e.g., Donaldson and Pendakur, 2004). Using the Euler equations from Section 4.3, our model predicts more savings than would be predicted if we take the OECD-modified equivalence scale $\eta = 1.5$. Hence, households in the Netherlands have relatively high economies of scale, implying they can save more.

The additional tax for singles to balance the government budget is $\widehat{\text{Tr}} = \widehat{\text{Tr}}_{\text{SS}} + \widehat{\text{Tr}}_{\text{LTC}} = 783.58 - 467.26 =$ €316.32 (for couples, this is double the amount). This consists of an additional tax to finance the first pillar pension ($\widehat{\text{Tr}}_{\text{SS}} > 0$) and a subsidy to finance the LTC system ($\widehat{\text{Tr}}_{\text{LTC}} < 0$). Given the low amounts we are talking about, we can think of these transfers reflecting measurement error due to calibration of the first-stage parameters.

6 Welfare Gain due to lower LTC use and Mortality

In this section, we closely follow De Nardi et al. (2023) and use the estimated life-cycle model to quantify the welfare gain arising from socioeconomic differences in LTC use and mortality (cf. Table 1). In the first step, we compute the monetary gain for any lifetime income quintile by counterfactually assigning them the health risks of the lowest lifetime income quintile. Besides, we evaluate the total welfare gain with a Willingness-To-Accept (WTA) metric that includes a non-monetary gain linked to reaching higher utility: the compensated consumption equivalence. As a final step, we utilize the unique feature of life-cycle models that allows us to quantify the extent to which saving for a bequest and the existence of LTC co-payments contribute to the observed WTAs.

6.1 Counterfactual Analyses

At age 65, households draw an LTC use and mortality risk profile that depends on their lifetime income quintile, denote this baseline scenario by BS. We also have a counterfactual scenario, denoted by CF, where each household draws the health risks of the lowest lifetime income quintile, so health risks are homogenous. The counterfactual implies that higher lifetime income quintiles live shorter, so have lower lifetime retirement income, and have higher lifetime LTC use, so have higher lifetime co-payments for LTC. Furthermore, lifetime co-payments will be different under the counterfactual due to the endogeneity of assets. Lastly, lifetime government-balancing transfers will be different due to lower longevity and because we will re-calibrate $\widehat{\mathrm{Tr}}_{\mathrm{SS}}$ and $\widehat{\mathrm{Tr}}_{\mathrm{LTC}}$ to also match the government budget under the counterfactual.

We compute the net present value of retirement income net of co-payments and government-balancing transfers and take the difference between baseline and counterfactual scenarios as the monetary gain from heterogeneous health risks. We do this for each lifetime income quintile separately. In concordance with LTC use and mortality risk starting, we measure the net present value at age 65. For the two cases, denote with $y^{BS}(\aleph_t)$ and $y^{CF}(\aleph_t)$ the net incomes for a household aged $t \geq 65$ with state vector $\aleph_t$. Denote $\mathbb{E}_{65}(y^{BS}(\aleph_t))$ and $\mathbb{E}_{65}(y^{CF}(\aleph_t))$ their expected values measured when the household is 65. These expectations are unconditional upon survival after age $t \geq 65$ and thus include differential mortality. The difference $\mathbb{E}_{65}(y^{BS}(\aleph_t)) - \mathbb{E}_{65}(y^{CF}(\aleph_t))$ is the contribution of age t to the monetary gain, and the expected lifetime income gain is the sum of the age-specific gains:

$$\sum_{t=65}^{100} \frac{\mathbb{E}_{65}(y^{BS}(\aleph_t)) - \mathbb{E}_{65}(y^{CF}(\aleph_t))}{(1+r)^{t-65}}, \tag{Monetary gain}$$

where we deflate the income stream to age 65 with an interest rate of $r = 0.02$. Apart from this level estimate, we will decompose the monetary gain into parts stemming from pension income, LTC co-payments, and government transfers.

Because our counterfactual affects consumption decisions, and bequest decisions, and the utility of life expectancy, we follow De Nardi et al. (2023) and adopt the compensated consumption equivalence λ_c as a measure of the welfare gain. This measure is the minimum percentage points increase in counterfactual consumption that a household requires to prefer (accept) the 'worse' counterfactual over the baseline case, hence a Willingness-To-Accept (WTA).

Formally, the expected lifetime utility at age 65 in the baseline scenario, value function V_{65}^{BS}, is defined as:

$$\begin{aligned} V_{65}^{BS} &:= \sum_{t=65}^{100} \hat{\beta}^{t-65} \cdot \left\{ \mathbb{E}_{65}\left(u(c_t^{BS}(\aleph_t))\right) + \hat{\beta} \cdot \mathbb{E}_{65}\left(\mathcal{B}(a_{t+1}^{BS}(\aleph_t))\right)\right\} \\ &= \sum_{t=65}^{100} \hat{\beta}^{t-65} \cdot \left\{ \mathbb{E}_{65}\left((1+\mathbb{1}(\aleph_t)) \cdot \frac{\left(\frac{c_t^{BS}(\aleph_t)}{\hat{\eta}(\aleph_t)}\right)^{1-\sigma}}{1-\sigma} + \bar{b}\right) + \hat{\beta} \cdot \mathbb{E}_{65}\left(\hat{c_a}^{-\sigma} \cdot a_{t+1}^{BS}(\aleph_t)\right)\right\}, \end{aligned}$$

which is the sum of expected lifetime utility from consumption and bequeathing. c_t^{BS} and a_{t+1}^{BS} are optimal consumption and a bequest at age t and $t+1$ for a household endowed with state vector $\aleph_t$. Note bequest utility is linear in assets because we estimate $\hat{\phi} = 1$.

Similarly, we determine the optimal consumption c_t^{CF}, bequests a_{t+1}^{CF}, and value function V^{CF} for the counterfactual case:

$$V^{CF}(\lambda_c) := \sum_{t=65}^{100} \hat{\beta}^{t-65} \cdot \left\{ \mathbb{E}_{65}\left((1+\mathbb{1}(\aleph_t)) \cdot \frac{\left(\frac{(1+\lambda_c)\cdot c_t^{CF}(\aleph_t)}{\hat{\eta}(\aleph_t)}\right)^{1-\sigma}}{1-\sigma} + \bar{b}\right) + \hat{\beta} \cdot \mathbb{E}_{65}\left(\hat{c_a}^{-\sigma} \cdot a_{t+1}^{CF}(\aleph_t)\right)\right\}.$$

We find WTA λ_c by solving: $V^{CF}(\lambda_c) = V_{65}^{BS}$. Without compensating ($\lambda_c = 0$), we expect less lifetime utility in the counterfactual scenario: $V^{CF}(0) < V_{65}^{BS}$. Because the utility is increasing in consumption, we have $\frac{\partial V^{CF}(\lambda_c)}{\partial \lambda_c} > 0$, and thus require $\lambda_c > 0$ to have $V^{CF}(\lambda_c) = V_{65}^{BS}$. $\lambda_c > 0$ represents the welfare gain: the closer this number is to zero, the smaller the welfare gain for the lifetime income quintile. Due to different deflation, we cannot directly compare λ_c to the monetary gain within a lifetime income quintile: monetary gains are obtained by using discount factor $\frac{1}{1+r}$, while λ_c is obtained by using discount factor $\hat{\beta} < \frac{1}{1+r}$.

In a final step, we look at the impact of LTC co-payments and saving for a bequest on the WTA. To this end, we one-by-one remove LTC co-payments and saving for a bequest ($\phi = 0$) for the baseline case and recompute optimal c_t^{BS} and a_{t+1}^{BS}, so V_{65}^{BS}. For the counterfactual, we keep c_t^{CF} and a_{t+1}^{CF} fixed and find λ_c that solves $V^{CF}(\lambda_c) = V_{65}^{BS}$.

6.2 Results

The first two lines in Table 5 show the average gain in lifetime income if health risks differ by lifetime income quintile, i.e., higher lifetime income quintiles use less case and live longer. The pecuniary gain per income group reveals a gradient favoring higher lifetime income quintiles. However, this result is incomplete because higher lifetime income quintiles by construction have higher absolute gains due to higher yearly income. To back out these level effects, the second row shows the gain relative to group-specific lifetime income under the counterfactual. The result confirms the gradient: the gain is −0.2% for the bottom income group and 11.0% for the top income group, a difference of 11.2pp..

When discussing welfare gains, we prefer the first-differenced estimate of 11.2pp., which accounts for the fact that bottom income groups, despite unchanged health risks, still loss or gain welfare under the counterfactual. The bottom income group namely loses 0.2% of lifetime income due a re-calibration of $\widehat{\mathrm{Tr}}_{\mathrm{SS}}$ and $\widehat{\mathrm{Tr}}_{\mathrm{LTC}}$. This gain is not directly linked to differences in health and is common to all income groups, and therefore we prefer the first-differenced estimate of 11.2pp..

With shares over 90%, we see that pension income is the largest contributor to the pecuniary gain for a lifetime income quintile. The role of LTC co-payments is non-negligible for the highest lifetime income quintile and explains 10.6% (€10,878) of their pecuniary gain. As a side-remark, the co-payments make up a small yet negative share for the second lifetime income quintile because their baseline LTC use is *higher* than under the counterfactual (see Table 1).

Table 5: Monetary and Welfare Gains Due To Socioeconomic Differences in LTC Use and Mortality: Levels and Decomposition

Lifetime income quintile	Bottom	Second	Third	Fourth	Top	ΔTop-Bottom
Monetary gain (€)	-917	14,075	26,864	49,812	102,474	103,391
Monetary gain[1] (%)	-0.2	3.2	5.2	8.0	11.0	11.2
Contribution to monetary gain[2] (%)						
Pension income	0.0	112.2	103.3	96.8	91.3	-
Co-payments	0.1	-2.8	2.3	6.8	10.6	-
Government transfers Tr_x	99.9	-9.4	-5.6	-3.5	-1.9	-
Willingness-To-Accept: $\lambda_c \times 100\%$	-0.2	0.8	2.9	7.9	23.2	23.4
No bequests ($\phi = 0$)	-0.5	0.1	1.7	4.9	0.7	1.2
No co-payments	2.4	3.6	5.6	10.4	24.2	21.8
No co-payments and bequests	2.3	3.4	5.1	8.7	3.3	1.0

[1] Expressed as a percentage of counterfactual lifetime income after age 65
[2] Gain of the particular income source in €s as a share of the monetary gain in €s (first row)

The WTA confirms higher welfare gains for higher lifetime income quintiles, but what explains the gap of 23.4pp.? If we assume away saving for a bequest, the gap in welfare gain (WTA) between the top and bottom lifetime income quintiles shrinks from 23.4pp. to 1.2pp.. Hence, higher lifetime income quintiles benefit less from higher longevity and lower LTC use if they cannot save for a bequest. Their welfare gain dropped from 23.2% to 0.7% because they will not spend the additional lifetime income on their otherwise highly-valued bequests. On the contrary, lower lifetime income quintiles experience a much smaller drop in welfare gain because they value leaving bequests –luxury goods– much less. As a result, the difference in welfare gain between the top and bottom lifetime income quintile shrinks tremendously.

On the other hand, the gap between the top and bottom income groups remains a considerable 21.8pp. when we leave out LTC co-payments. Differences in co-payment duration are thus less important than a bequest to explain the excess welfare gain of the top lifetime income quintile. The gap remains large because LTC co-payments are a relatively small share of lifetime income gains (Row 4, Table 5). Moreover, higher lifetime income quintiles still receive the additional retirement income, which they can spend on –for them valuable– bequests.

Note that if we abolish LTC co-payments in the baseline scenario, any group experiences a welfare gain, which is good from a social planner's perspective. While their risks are not altered, the bottom lifetime income quintile has a welfare gain 2.4% because LTC co-payments are replaced by a higher transfer (tax) $\widehat{\mathrm{Tr}}_{\mathrm{LTC}}$ that is paid *unconditionally* upon LTC use. Lower socioeconomic groups can spend the otherwise co-paid resources on consumption, while the same is true for the higher socioeconomic groups who can additionally spend it on for them valuable bequests.

If we simultaneously assume away saving for bequests and LTC co-payments, we find welfare gains that are in between the cases of singling out only one of the two channels. In line with the evidence above, for lower lifetime income quintiles, the gain is closest to the case of singling out LTC co-payments only. In comparison, for higher lifetime income quintiles, the case is closer to singling out bequests only, as these are more valuable for them.

7 Discussion and conclusion

We evaluate the welfare gain that Dutch households with higher lifetime income experience due to using less long-term care (LTC) and living longer. To this end, we estimated a life-cycle model on singles and couples' consumption and saving behavior, including idiosyncratic risks on income, LTC use, and mortality. We calibrated the model to match Dutch administrative data on asset holdings from 2006-2014. Using the estimated model, we conducted three counterfactual experiments to quantify and explore possible channels of the welfare gain: (1) assign each household the LTC use and mortality risk of the bottom lifetime income quintile, (2) additionally remove a preference for bequest saving, and (3) replace co-payments for LTC with a fixed tax that is paid unconditionally upon using LTC.

Our findings highlight a sizeable excess welfare gain of 23.4pp. higher consumption for the highest lifetime income quintile if their health follows the true process rather than the counterfactual one. The large welfare gain for the top lifetime income quintile can almost exclusively be attributed to their preference for leaving bequests: the welfare gain reduces to 1.2pp. if households would not hold a preference for bequest saving. Our ranking exercise shows that LTC co-payments are less important when explaining the excess welfare gain: the gap remains 21.8pp..

The estimated welfare effects should be interpreted as a lower bound estimate because we calibrate the utility of remaining life-expectancy $\bar{b}$ at the lower end of possible values. This seems a sensible choice as Hall and Jones (2007) show that lower values of $\bar{b}$ better match healtcare expenditures in the U.S.. Yet it must be said that, in line with Hall and Jones (2007) and our own computations (not shown), the estimated welfare effects are sensitive to higher choices of $\bar{b}$.

In line with our findings, earlier work emphasized that modeling the bequest saving motive is crucial for understanding the asset holdings (welfare) of the income- and asset-rich (De Nardi et al., 2010). However, earlier work is primarily conducted in the U.S., where public LTC provision is less generous: savings data alone need not separately identify precautionary and bequest motives because wealthier households simultaneously save assets for both uses (Dynan et al., 2004). Our study is one of the first attempts to estimate the bequest saving motive in a country where precautionary saving is less important, and thus saving data alone could suffice. From our findings, we conclude that the estimated preference for bequest saving seems consistent across countries and estimation strategies.

For policy design, we can conclude that higher taxes on bequests could be a way to (partially) introduce more 'actuarial fairness' into the system of old-age social insurances. Also, lower lifetime income quintiles do not necessarily have a worse deal from abolishing co-payments for LTC. Yet, it must be said that such policy intervention will also not close the gap with the highest lifetime income quintile tremendously. Furthermore, co-payments works out differently in another country because co-payments are relatively low in the Netherlands (OECD, 2023). Along the spectrum of possible policy interventions, having social security benefits tailored to the career length is another way to increase actuarial fairness because the working life of shorter-living (lower) lifetime income quintiles usually starts at younger ages. While our findings opt for those kinds of policies, we keep the quantitative importance of these alternative policy proposals and their interaction with heterogenous mortality *and* LTC use as a fruitful area for future work (see Bagchi (2019) for an example involving differential mortality only).

Looking at policy reforms would make the working age stage more salient than it is now. In the current setup of studying health differences, we primarily use this life-cycle stage to close the government budget. Alternatively, focusing on bigger reforms to the social insurance system increases the relevance of precautionary saving, which already happens during working age. Furthermore, the comparison of alternative systems implicitly compares different contributory schemes, and these contributions mostly occur during working age.

Our analysis could be extended in another number of directions in future work. First, when assessing different retirement policies, we might have to extend the life-cycle model with endogenous health and labor supply decisions and health-dependent utility (cf. French, 2005; Finkelstein et al., 2009). In our specification, we pursue parsimony and thus treat health as exogenous and do not model labor supply explicitly. That does not mean we entirely ignore these variables; the income risk reflects them and, therefore lifetime income status at age 65. Second, future research could seek to estimate the effects of changing the LTC insurance system besides the co-payments. In doing so, we can assess whether our studied welfare gains are larger in a system with exclusively private LTC insurance (saving) or a mix of public and private LTC insurance. Lastly, future research can include other behavioral frictions that likely matter for evaluating of the impact of bequests. The typical frictions that one can think of are taxes on bequests, taxes on capital gains, and trade-offs between leaving bequests and inter-vivos transfers (gifts), which are not part of our model.

References

Altonji, J. G. and Segal, L. M. (1996). Small-Sample Bias in GMM Estimation of Covariance Structures. *Journal of Business & Economic Statistics*, 14(3):353–366.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.

Ameriks, J., Caplin, A., Laufer, S., and van Nieuwerburgh, S. (2011). The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives. *The Journal of Finance*, 66(2):519–561.

Arellano, M. (2003). *Panel Data Econometrics*. Oxford University Press, Oxford (United Kingdom).

Auerbach, A. J., Charles, K. K., Coile, C. C., Gale, W., Goldman, D., Lee, R., Lucas, C. M., Orszag, P. R., Sheiner, L. M., Tysinger, B., Weil, D. N., Wolfers, J., and Wong, R. (2017). How the growing gap in life expectancy may affect retirement benefits and reforms. NBER Working Papers no. 23329, National Bureau of Economic Research.

Bagchi, S. (2019). Differential mortality and the progressivity of social security. *Journal of Public Economics*, 177:104044.

Bakx, P., Van Doorslaer, E., and Wouterse, B. (2023). Long-term care in the Netherlands. NBER Working Papers no. 31823, National Bureau of Economic Research.

Barczyk, D. and Kredler, M. (2018). Evaluating long-term-care policy options, taking the family seriously. *The Review of Economic Studies*, 85(2):766–809.

Bernheim, B. D. (1987). The economic effects of social security: toward a reconciliation of theory and measurement. *Journal of Public Economics*, 33(3):273–304.

Bhattacharya, J. and Lakdawalla, D. (2006). Does Medicare benefit the poor? *Journal of Public Economics*, 90(1):277–292.

Blundell, R., Graber, M., and Mogstad, M. (2015). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics*, 127:58–73.

Bosworth, B., Burtless, G., and Zhang, K. (2016). Later retirement, inequality in old age, and the growing gap in longevity between rich and poor. *Economic Studies at Brookings*.

Braun, R. A., Kopecky, K. A., and Koreshkova, T. (2017). Old, sick, alone, and poor: a welfare analysis of old-age social insurance programmes. *The Review of Economic Studies*, 84(2):580–612.

Browning, M., Chiappori, P.-A., and Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, 80(4 (285)):1267–1303.

Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics*, 21(3):339–353.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55.

Cherchye, L., Demuynck, T., De Rock, B., Kovaleva, M., Minne, G., Perea, M. D. S., and Vermeulen, F. (2023). Poor and wealthy hand-to-mouth households in Belgium. *Review of Economics of the Household*.

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., and Cutler, D. (2016). The association between income and life expectancy in the United States, 2001-2014. *JAMA*, 315(16):1750–1766.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy*, 118(1):39–75.

De Nardi, M., French, E., and Jones, J. B. (2016). Medicaid insurance in old age. *American Economic Review*, 106(11):3480–3520.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021). Why do couples and singles save during retirement? Household heterogeneity and its aggregate implications. NBER Working Papers no. 28828, National Bureau of Economic Research.

De Nardi, M., Pashchenko, S., and Porapakkarm, P. (2023). The lifetime costs of bad health. *Review of Economic Studies (forthcoming)*.

Deaton, A. (2002). Policy implications of the gradient of health and wealth. *Health Affairs*, 21(2):13–30.

Donaldson, D. and Pendakur, K. (2004). Equivalent-expenditure functions and expenditure-dependent equivalence scales. *Journal of Public Economics*, 88(1):175–208.

Dynan, K. E., Skinner, J., and Zeldes, S. P. (2004). Do the rich save more? *Journal of Political Economy*, 112(2):397–444.

Fehr, H., Kallweit, M., and Kindermann, F. (2013). Should pensions be progressive? *European Economic Review*, 63:94–116.

Finkelstein, A., Luttmer, E. F. P., and Notowidigdo, M. J. (2009). Approaches to estimating the health state dependence of the utility function. *American Economic Review*, 99(2):116–121.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

Goda, G., Shoven, J., and Slavov, S. (2011a). Differential Mortality by Income and Social Security Progressivity. In *Explorations in the Economics of Aging*, volume 5, pages 189 – 204. University of Chicago Press, Chicago.

Goda, G. S., Golberstein, E., and Grabowski, D. C. (2011b). Income and the utilization of long-term care services: evidence from the social security benefit notch. *Journal of Health Economics*, 30(4):719–729.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Groneck, M. and Wallenius, J. (2021). It sucks to be single! Marital status and redistribution of social security. *The Economic Journal*, 131(633):327–371.

Hall, R. E. and Jones, C. I. (2007). The value of life and the rise in health spending*. *The Quarterly Journal of Economics*, 122(1):39–72.

Heathcote, J., Violante, G., and Storesletten, K. (2020). Presidential address 2019: how should tax progressivity respond to rising income inequality? 18(6):2715–2754.

Jones, J. B., De Nardi, M., French, E., McGee, R., and Kirschner, J. (2018). The lifetime medical spending of retirees. NBER Working Papers no. 24599, National Bureau of Economic Research.

Jones, J. B. and Li, Y. (2023). Social security reform with heterogeneous mortality. *Review of Economic Dynamics*, 48:320–344.

Karahan, F. and Ozkan, S. (2013). On the persistence of income shocks over the life cycle: evidence, theory, and implications. *Review of Economic Dynamics*, 16(3):452–476.

Knoef, M., Been, J., Alessie, R., Caminada, K., Goudswaard, K., and Kalwij, A. (2016). Measuring retirement savings adequacy: developing a multi-pillar approach in the Netherlands. *Journal of Pension Economics & Finance*, 15(1):55–89.

Knoef, M., Been, J., Rhuggenaath, J., Goudswaard, C., and Caminada, K. (2017). De toereikendheid van pensioenopbouw na de crisis en pensioenhervormingen. Netspar Design Paper no. 68, Netspar.

Ko, A. (2022). An equilibrium analysis of the long-term care insurance market. *The Review of Economic Studies*, 89(4):1993–2025.

Kopecky, K. A. and Suen, R. M. H. (2010). Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13(3):701–714.

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–2550.

McClellan, M. and Skinner, J. (2006). The incidence of Medicare. *Journal of Public Economics*, 90(1):257–276.

Mincer, J. A. (1974). *Schooling, experience, and earnings.* Columbia University Press.

Nakajima, M. and Telyukova, I. (2023). Medical expenses and saving in retirement: the case of U.S. and Sweden. *American Economic Journal: Macroeconomics (conditionally accepted).*

OECD (2023). *Health at a glance 2023: OECD indicators.* OECD, Paris.

Paz-Pardo, G. and Galves, J. (2023). Richer earnings dynamics, consumption and portfolio choice over the life-cycle. Working Paper.

Poterba, J. M. (2014). Retirement Security in an Aging Population. *American Economic Review*, 104(5):1–30.

Powell, J. L. (1994). Estimation of semiparametric models. In Engle, R. F. and McFadden, D. L., editors, *Handbook of Econometrics*, volume 4, pages 2443–2521. Elsevier.

Rodrigues, R., Ilinca, S., and Schmidt, A. E. (2018). Income-rich and wealth-poor? The impact of measures of socio-economic status in the analysis of the distribution of long-term care use among older people. *Health Economics*, 27(3):637–646.

Rouwenhorst, G. (1995). Asset pricing implications of equilibrium business cycle models. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, New Jersey, Princeton.

Smith, J. (2007). The impact of socioeconomic status on health over the life-course. *Journal of Human Resources*, 42(4).

St-Amour, P. (2022). Valuing life over the life cycle. Working paper.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633.

Tenand, M., Bakx, P., and van Doorslaer, E. (2020a). Equal long-term care for equal needs with universal and comprehensive coverage? An assessment using Dutch administrative data. *Health Economics*, 29(4):435–451.

Tenand, M., Wouterse, B., Bakx, P., and Bom, J. (2020b). Do co-payments affect nursing home entry? An assessment based on Dutch administrative data. Netspar Design Paper no. 139.

van den Berg, G. (2001). Handbook of Econometrics. In Heckman, J. and Leamer, E., editors, *Handbook of Econometrics*, volume 5, pages 3381–3460.

van der Vaart, J., Alessie, R., Groneck, M., and van Ooijen, R. (2023). Combining insurance against old-age risks to accommodate socioeconomic differences in long-term care use and mortality. Working paper.

van der Vaart, J., Alessie, R., and van Ooijen, R. (2020). Economic consequences of widowhood: evidence from a survivor's benefits reform in the Netherlands. Netspar Design Paper no. 160.

van Ooijen, R., Alessie, R., and Kalwij, A. (2015). Saving behavior and portfolio choice after retirement. *De Economist*, 163(3):353–404.

Wouterse, B., Hussem, A., and Wong, A. (2021). The risk protection and redistribution effects of long-term care co-payments. *Journal of Risk and Insurance*, 89(1):161–186.

Appendices

A Life-cycle Model

A.1 Government Budget Constraint

The government collects the taxes and co-payments to finance expenditures on the first pillar pension and LTC provision. Yet, government revenues and spending are not guaranteed to be balanced in the model. To let the government break even, we assume additional fixed transfers of Tr_{SS} and Tr_{LTC}(a tax or subsidy) in each age period. For a household of a given age, the government expenditures on LTC are:

$$\text{LTC}(h_t^m, h_t^f) = \begin{cases} 2 \cdot \text{LTC}_{\text{cost}} & \text{if } \; h_t^m = 2 \; \text{ and } \; h_t^f = 2, \\ \text{LTC}_{\text{cost}} & \text{if } \; h_t^m = 2 \; \text{ or } \; h_t^f = 2, \\ 0 & \text{elsewhere,} \end{cases}$$

where LTC_{cost} = €58, 500 is the cost of an individual stay in a public institution for a year.

Similarly, the government pays first pillar pension:

$$\text{SS}(t, h_t^m, h_t^f) = \begin{cases} 2 \cdot w & \text{if } \; t \geq 65, \; h_t^m \neq 3 \; \text{ and } \; h_t^f \neq 3, \\ 1.4 \cdot w & \text{if } \; t \geq 65, \; h_t^m = 3 \; \text{ or } \; h_t^f = 3, \\ 0 & \text{elsewhere.} \end{cases}$$

These are the expenditures per household and conditional upon age t and health statuses h_t^m and h_t^f. Total, i.e., unconditional, government expenditures GE_{LTC} and GE_{SS} are the expenditures per household weighted by the steady-state distribution on household types $f(\aleph)$, with $\aleph = \aleph^W \cup \aleph^R = (a_t, \theta, \eta_t, \epsilon_t, \text{DB}_t, f_t, h_t^m, h_t^f, t)'$. Then:

$$\text{GE}_{SS} = \sigma_1 \cdot \int_{\aleph} f(\aleph) \cdot \text{SS}(\aleph) d\aleph, \quad \text{and} \quad \text{GE}_{LTC} = \sigma_2 \cdot \int_{\aleph} f(\aleph) \cdot \text{LTC}(\aleph) d\aleph,$$

where σ_1 and σ_2 reflect the share of government expenditures financed through dedicated taxes and co-payments. The rest is financed with general taxes and not of interest when balancing the government budget.[13]

To finance these benefits, the government obtains revenue from taxes and co-payments: $\tau_{\mathrm{SS}}(\cdot)$, $\tau_{\mathrm{L}}(\cdot)$, and $m(\cdot)$. Also, there is an additional balancing transfer Tr_x with $x \in (SS, L)$. The transfer is defined as follows:

$$\mathrm{Tr}_x(f) = \begin{cases} 2 \cdot \mathrm{Tr}_x & \text{if } f = \text{ couple} \\ \mathrm{Tr}_x & \text{if } f = \text{ single woman or single man,} \end{cases}$$

and is thus twice as large for couples than for singles.

Government revenues, GR_x, are given by:

$$\mathrm{GR}_{SS}(\mathrm{Tr}_{SS}) = \int_{\aleph} f(\aleph) \cdot (\tau_{\mathrm{SS}}(\aleph) + \mathrm{Tr}_{SS}(\aleph))\, \mathrm{d}\aleph \quad \text{and}$$
$$\mathrm{GR}_{LTC}(\mathrm{Tr}_{LTC}) = \int_{\aleph} f(\aleph) \cdot (\tau_{\mathrm{L}}(\aleph) + m(\aleph) + \mathrm{Tr}_{LTC}(\aleph))\, \mathrm{d}\aleph,$$

which consist of the sum of taxes, co-payments for LTC, and the additional tax (subsidy) that balances the government budget constraint.

The government sets the transfer levels Tr_x according to:

$$\mathrm{GE}_x = \mathrm{GR}_x(\mathrm{Tr}_x),$$

which can be tax or subsidy, depending on whether there is a deficit or a surplus. Appendix A.4 explains how we compute these transfers numerically.

[13]We take the values from 2010: $\sigma_1 = 0.664$ and $\sigma_2 = 0.640$, which we computed using aggregate expenditures and revenues reported on: https://www.cbs.nl/nl-nl/nieuws/2019/37/inkomsten-uit-sociale-premies-6-1-miljard-hoger-in-2018, [August 7, 2023] and https:/opendata.cbs.nl/statline/CBS/nl/dataset/84121NED/table?ts=1564565763409, [August 7, 2023].

A.2 Closed-form Solution for Policy Function Iteration

We elaborate here on how the households determine their consumption policy functions. We use the Bellmann maximization principle, which recursively solves the household optimization problem from the last to the first life-cycle period. The objective function is the value function in this case. A general form of the value function in any state $\aleph = \aleph^W \cup \aleph^R$ is given by:

$$
\begin{aligned}
\mathrm{V}(\aleph;\ h_t^m = i, h_t^f = j) = \max_{c_t, a_{t+1}}\ & \mathrm{u}^f\left(c_t\right) + \beta \cdot \Big(\big(1 - \pi_{3,3}^{i,j}(t,I)\big) \cdot \mathbb{E}[\mathrm{V}(\aleph^+)|\aleph] + \pi_{3,3}^{i,j}(t,I) \cdot \mathcal{B}(a_{t+1}) \Big) \\
\text{s.t.}\quad & a_{t+1} = R \cdot a_t + y_t - \tau_{SS} - \tau_L - \tau_G - m_t + \mathrm{Tr}_{SS}(\cdot) + \mathrm{Tr}_{LTC}(\cdot) - c_t \\
& a_{t+1} \geq 0,
\end{aligned} \tag{3}
$$

where $(i,j) \in \{1,2,3\}$. The Lagrangian corresponding to (3) reads as:

$$
\begin{aligned}
\max_{c_t, a_{t+1}, \lambda} \mathcal{L}(\cdot) = \ & \mathrm{u}^f\left(c_t\right) + \beta \cdot \Big(\big(1 - \pi_{3,3}^{i,j}(t,I)\big) \cdot \mathbb{E}[\mathrm{V}(\aleph^+)|\aleph] + \pi_{3,3}^{i,j}(t,I) \cdot \mathcal{B}(a_{t+1}) \Big) \\
& + \lambda \cdot \{R \cdot a_t + y_t - \tau_{SS} - \tau_L - \tau_G - m_t + \mathrm{Tr}_{SS}(\cdot) + \mathrm{Tr}_{LTC}(\cdot) - c_t - a_{t+1}\},
\end{aligned} \tag{4}
$$

which has the following first-order constraints:

$$
\frac{\partial \mathcal{L}(\cdot)}{\partial c_t} := u_{c^t}^f - \lambda = 0 \tag{5}
$$

$$
\frac{\partial \mathcal{L}(\cdot)}{\partial a_{t+1}} := \beta \cdot \Big((1 - \pi_{3,3}^{i,j}(t,I) \cdot \mathbb{E}[\mathrm{V}_{a_{t+1}}(\aleph^+)|\aleph] + \pi_{3,3}^{i,j}(t,I) \cdot \mathcal{B}_{a_{t+1}}(a_{t+1}) \Big) - \lambda = 0 \tag{6}
$$

$$
\frac{\partial \mathcal{L}(\cdot)}{\partial \lambda} := R \cdot a_t + y_t - \tau_{SS} - \tau_L - \tau_G - m_t + \mathrm{Tr}_{SS}(\cdot) + \mathrm{Tr}_{LTC}(\cdot) - c_t - a_{t+1} = 0. \tag{7}
$$

Note that $\mathrm{V}(\aleph)$ in (3) is an optimum, and so is the Lagrangian in (4) when analyzed in $c_t(\aleph)$, $a_{t+1}(\aleph)$, and $\lambda(\aleph)$. As a consequence, we can apply the envelope theorem:

$$
\mathrm{V}_{a_t}(\aleph) = \frac{\partial \mathrm{V}(\aleph)}{\partial a_t} = \frac{\partial \mathcal{L}(\cdot)}{\partial a_t}\bigg|_{c_t(\aleph), a_{t+1}(\aleph), \lambda(\aleph)} = \lambda(\aleph) \cdot R,
$$

which has to hold in the next period as well:

$$\mathrm{V}_{a_{t+1}}(\aleph^+) = \frac{\partial \mathrm{V}(\aleph^+)}{\partial a_{t+1}} = \frac{\partial \mathcal{L}(\cdot)}{\partial a_{t+1}}\bigg|_{c_{t+1}(\aleph^+), a_{t+2}(\aleph^+), \lambda(\aleph^+)} = \lambda(\aleph^+) \cdot R, \tag{8}$$

where $\aleph^+$ is the state vector in the next period. Furthermore, (5) holds optimally in the future:

$$u^{f^+}_{c_{t+1}}(c_{t+1}(\aleph^+)) = \lambda(\aleph^+). \tag{9}$$

Combining (8) and (9) yields:

$$\mathrm{V}_{a_{t+1}}(\aleph^+) = u^{f^+}_{c_{t+1}}(c_{t+1}(\aleph^+)) \cdot R \tag{10}$$

Using (10), we build the Euler equation that describes the evolution of consumption and assets over time. We combine (10) with (5) and (6), while (7) simultaneously holds (together with the non-negativity constraint of assets). The Euler equation on consumption and bequests (assets) is:

$$\begin{aligned} u^f_{c_t}(c_t(\aleph)) &= \beta \cdot \big(\, (1 - \pi^{i,j}_{3,3}(t, I)) \cdot R \cdot \mathbb{E}[u^{f^+}_{c_{t+1}}(c_{t+1}(\aleph^+)) \, | \aleph] \\ &+ \pi^{i,j}_{3,3}(t, I) \cdot \mathcal{B}_{a_{t+1}}(a_{t+1}(\aleph))\big) \qquad (11) \\ \text{with:} \quad a_{t+1}(\aleph) &= R \cdot a_t + y_t - \tau_{SS} - \tau_L - \tau_G - m_t + \mathrm{Tr}_{SS}(\cdot) + \mathrm{Tr}_{LTC}(\cdot) - c_t(\aleph) \; \geq \; 0 \end{aligned}$$

This system can be recursively solved if we know the solution for the last period.

A.3 Terminal Period Solution

We now solve the dynamic program problem for the terminal (last) period $t = T$. Note that the household will not be around in the next period ($\pi^{i,j}_{3,3}(T, I) = 1$) but can bequeath, where $(i, j) \in \{1, 2, 3\}$. The terminal period solution of (11) in state $\aleph$ reduces to:

$$\begin{aligned} u^f_{c_T}(c_T(\aleph)) &= \beta \cdot \mathcal{B}_{a_{T+1}}(a_{T+1}(\aleph)) \\ a_{T+1}(\aleph) &= R \cdot a_T + y_T - \tau_{SS} - \tau_L - \tau_G - m_T + \mathrm{Tr}_{SS}(\cdot) + \mathrm{Tr}_{LTC}(\cdot) - c_t(\aleph) \\ &= \mu - c_T(\aleph) \; \geq \; 0 \qquad (12) \end{aligned}$$

where μ is the total wealth holding at age T that is split over consumption and a bequest.

To solve the system, we have to consider three cases: $\phi = 0$ (no bequest), $\phi \in (0,1)$ (some wealth above threshold c_a is bequeathed), and $\phi = 1$ (all wealth above threshold c_a is bequeathed). The marginal utility of leaving a bequest is:

$$\mathcal{B}_{a_{T+1}}(a_{T+1}) = \begin{cases} 0 & \text{if } \phi = 0 \\ \frac{\phi}{1-\phi}^{\sigma} \cdot \left(\frac{\phi}{1-\phi} \cdot c_a + a_{T+1}\right)^{-\sigma} & \text{if } \phi \in (0,1) \\ c_a^{-\sigma} & \text{if } \phi = 1. \end{cases}$$

Also, marginal utility from consumption depends on family structure:

$$u_{c_T}(c_T) = \begin{cases} c_T^{-\sigma} & \text{if } f_T = \text{ single man or woman} \\ 2 \cdot \left(\frac{1}{\eta}\right)^{1-\sigma} \cdot c_T^{-\sigma} & \text{if } f_T = \text{ couple.} \end{cases}$$

If $\phi = 0$, the Euler equation in (12) becomes:

$$\begin{aligned} u^f_{c_T}\left(c_T(\aleph)\right) &= \beta \cdot \mathcal{B}_{a_{J+1}}(a_{J+1}(\aleph)) \rightarrow \\ u^f_{c_T}\left(c_T(\aleph)\right) &> \beta \cdot 0 \rightarrow \\ c_T(\aleph, \mu) &= \mu, \end{aligned} \tag{13}$$

where the latter equality stems from the budget constraint in (12).

If $\phi = 1$, the Euler equation in (12) becomes:

$$\beta \cdot \mathcal{B}_{a_{J+1}}(a_{J+1}(\aleph)) = \begin{cases} c_T^{-\sigma} & \text{if } f_T = \text{ single man or woman} \\ 2 \cdot \left(\frac{1}{\eta}\right)^{1-\sigma} \cdot c_T^{-\sigma} & \text{if } f_T = \text{ couple.} \end{cases}$$

Solving for c_T gives:

$$c_T(\aleph, \mu) = \begin{cases} \min\left(\beta^{-\frac{1}{\sigma}} \cdot c_a,\ \mu\right) & \text{if } f_T = \text{ single man or woman} \\ \min\left(2^{\frac{1}{\sigma}} \cdot \left(\frac{1}{\eta}\right)^{\frac{1}{\sigma}-1} \cdot \beta^{-\frac{1}{\sigma}} \cdot c_a,\ \mu\right) & \text{if } f_T = \text{ couple.} \end{cases}$$

Similarly, we solve the Euler equation for $\phi \in (0,1)$ and get:

$$c_T(\aleph, \mu) = \min\left(\left(\frac{x_1(f_T)}{x_1(f_T)+x_2} \cdot x_1^{-1}(f_T) \cdot c_a + \frac{x_2}{x_1(f_T)+x_2} \cdot \mu\right)\ ,\ \mu\right) \tag{14}$$

with:

$$x_1^{-1}(f_T) = \begin{cases} \beta^{-\frac{1}{\sigma}} & \text{if } f_T = \text{ single man or woman} \\ 2^{\frac{1}{\sigma}} \cdot \left(\frac{1}{\eta}\right)^{\frac{1}{\sigma}-1} \cdot \beta^{-\frac{1}{\sigma}} & \text{if } f_T = \text{ couple,} \end{cases}$$

and $x_2 = \left(\frac{\phi}{1-\phi}\right)^{-1}$.

Note that the bequest size is $a_{T+1}(\aleph, \mu) = \max\left(\mu - c_T(\aleph, \mu)\ ,\ 0\right)$ in all cases.

A.4 Numerically Solving the Model

We first discretize the state space and then solve the model along the discrete space.

Discretizing the state space Consider the vector with state variables $\aleph = \aleph^W \cup \aleph^R = (a_t, \theta, \eta_t, \epsilon_t, \mathrm{DB}_t, f_t, h_t^m, h_t^f, t)'$. This vector contains continuous variables a_t, θ, η_t, DB_t, and ϵ_t. Solving the Euler equation for each value is computationally too demanding and we, therefore, discretize these variables while maintaining the core properties of their distribution.

We discretize labor productivity $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$ and the transitory income shock $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ into a five- and three-dimensional grid using Gauss–Hermite quadrature. We discretize the stochastic AR(1)-variable η_t into a time-independent three-state Markov process. We use the decomposition method by Rouwenhorst (1995), which preserves the unconditional mean, the unconditional variance, and the auto-correlation of the actual process. Kopecky and Suen (2010) describes the algorithm in detail. We discretize the second pillar pension benefit on a 12-dimensional exponential grid from 0 to $150,000$ (growth rate = 0.52).

Lastly, we discretize assets (a_t) over a grid $\hat{\mathcal{A}}$ from €0 to €1,000,000. The asset grid contains 100 values. To prevent oscillation of the model for asset levels near zero, we take an exponential grid, i.e., we take relatively more low than high values for assets a_t on the grid (growth rate = 0.05).

Solving the model We require the probability distribution of assets $a_{t+1}(\aleph)$ and consumption $c_t(\aleph)$ at any age t. Suppose all parameter values are known in the model. We apply policy function iteration to solve the model and then compute the probability distribution.

We start with the closed-form solution of the terminal period T provided in Appendix A.3. We hereafter numerically solve the Euler equation system (11) from period $T-1$ back to period 1 and calculate the resulting policy functions $c_t(\aleph)$ and $a_{t+1}(\aleph)$.

Next, we compute the distribution of households over the state space $\aleph$. To increase computational speed, we analytically compute the distribution rather than infer this from a simulation (see, e.g., Cagetti, 2003). Furthermore, directly computing the distribution prevents that in an agent-based simulation, it remains unknown for what number of households the model statistics converge.

We compute the state distribution at age t by updating the state distribution at time $t-1$. For this, we assume an initial state distribution at age $t = 25$. The initial household consists of a couple without using LTC. They draw labor productivity level θ from the discrete distribution. We take $a_0 = 0$, $\text{DB}_{25} = 0$, and $\eta_{24} = 0$, so the household initially has no assets, pension accruals, and income shock. This distribution is modified to create a distribution over the state space for age 26. Given the current state $\aleph$ at age 25, we know how many assets any household chooses to possess at age 26 and the conditional probability of ending up in a particular health and income state at age 26. This information (transition matrix) suffices to update the state distribution of $\aleph$ from age 25 to the distribution at age 26. We repeat this procedure until age $t = T = 100$.

These state distributions are also essential to compute the transfer Tr_x that would balance the government budget (see Appendix A.1), $x \in (SS, LTC)$. For each state, we know the cost of providing LTC and pension, the paid taxes, and co-payments. We can subsequently compute the expected government revenues and costs. We apply a bisection search to find the level Tr_x that exactly balances the revenues and cost.

B First-stage Estimates

B.1 Data and Estimation of the Health Processes

Socioeconomic differences in LTC use and mortality are the primary input in our analysis. To quantify them, we use longitudinal data on LTC use and mortality, a simulation model to compute complete life histories on LTC use and death, and a socioeconomic status measure to stratify the life histories. The data and estimation procedure of the health process closely follows van der Vaart et al. (2023), which we will summarize here.

We use unique registry data from Statistics Netherlands reporting an individual and household key, institutional care use, death, marital status, birth date, and gender for the Dutch population between 2006 and 2014. The data are unique due to their high frequency: the registers daily report whether an individual stays in an institution, i.e., a residential or nursing home, died, and has a partner, i.e., is married, has a partnership contract, or cohabits on a contractual basis. The high frequency of the data allows us to precisely model many short institutional care spells that occur. Furthermore, it will enable us to model the effect of marital status on LTC use and mortality precisely from the moment of marital dissolution onward.

We restrict the estimation of the health process to households whose members are both retired, i.e., aged 65 or older and have retirement income as their main income source. The age restriction seems natural as only 1.0% of the 65-year-olds in our sample uses institutional care. To save on the number of heterogeneous groups, and thus state space of the life-cycle model, we further restrict to individuals who are or were married at age 65. We observe 2,548,664 individuals and 1,487,109 households.

To construct a socioeconomic measure, we merge this data to household records on income – the sum of couple members' pre-tax income (incl. social transfers and pension income) – and financial assets (savings, stocks, and bonds). The socioeconomic status measure is the average sum of equivalized household income and annuitized financial assets (savings, stocks, and bonds), reflecting lifetime income. This comprehensive measure has the advantage that it considers that after retirement, some households have little income but many assets, e.g., former entrepreneurs (Knoef et al., 2016). We compute lifetime income quintiles $I \in \{1,2,3,4,5\}$ depending on quintiles of its distribution.[14]

To compute complete life histories on LTC use and death, we use the competing risk model from van der Vaart et al. (2023) that allows for socioeconomic dependencies in risks and explicitly accounts for the spouse as a potential informal care provider. We distinguish three individual states: not using public institutional care ($i = 1$), using public institutional care ($i = 2$), or death ($i = 3$). Home-based care use is not a separate state because its co-payments and, thus, redistributive effects are very limited in the Netherlands (Tenand et al., 2020b). For parsimony, marital status is modeled as a covariate, and not as a separate (sub-)state in the competing risk model. As a first step, we specify and estimate a proportional hazard model for the transition rate λ_{ij} of going from a given state i to state $j \neq i$ at age t (van den Berg, 2001):

$$\lambda_{ij}(t \mid \mathrm{mar}(t), G, I) = \exp\left(\gamma_{ij}(G,I) \cdot t + c_{ij}(G,I) + \beta_{ij}(G,I) \cdot \mathrm{mar}(t)\right) \tag{15}$$

where γ_{ij} is the age effect, c_{ij} is the effect of being single, and $c_{ij} + \beta_{ij}$ is the effect of having a partner ($\mathrm{mar}(t) = 0$: has no partner; $\mathrm{mar}(t) = 1$ has no partner). All coefficients are estimated conditional upon gender G and lifetime income quintile I.[15] We estimate the model following standard log-likelihood inference for duration models (van der Vaart et al., 2023)

Because we observe the relevant outcomes only between 2006 and 2014, we use the estimates of (15) to simulate complete life histories on LTC use, marital status, and mortality. We generate a survival probability and thus a random timing of the transition from i to j:

$$S_{ij}(t \mid \mathrm{mar}(t), G, I) \;=\; \mathbb{P}(T \geq t,\ j \mid \mathrm{mar}(t), G, I,\ i) \;=\; \exp\left(-\int_0^t \lambda_{ij}(\tau \mid \mathrm{mar}(\tau), G, I) d\tau\right) \tag{16}$$

The simulation starts at age 65 with 100,000 households, when both couple members are alive. Each individual can move to two possible destination states. Using (16), we draw

[14] An alternative would be to take the level of education, but the register on education is incomplete for older cohorts, implying we have to stick to the current data.

[15] See Appendix B.2 for the fit on LTC use and mortality. We also estimated a model including frailty, but this specification gave a worse fit on LTC use and mortality.

a transition time for each state. The minimum of the two transition times determines which actual transition occurs. We repeat this procedure for the successive states until both members died. While the simulation is finished for the couple member who dies first, we still have to simulate the life history of LTC use for the surviving partner after widowhood. We use (16) but take the dummy value $\mathrm{mar}(t) = 0$ instead of $\mathrm{mar}(t) = 1$. After this last spouse dies, we stop the simulation and have the complete –and dependent– life histories on LTC use and mortality for the two partners.

B.2 Goodness of Fit of Health Processes

Figure 2: Goodness of Fit of Survival Curves by Age

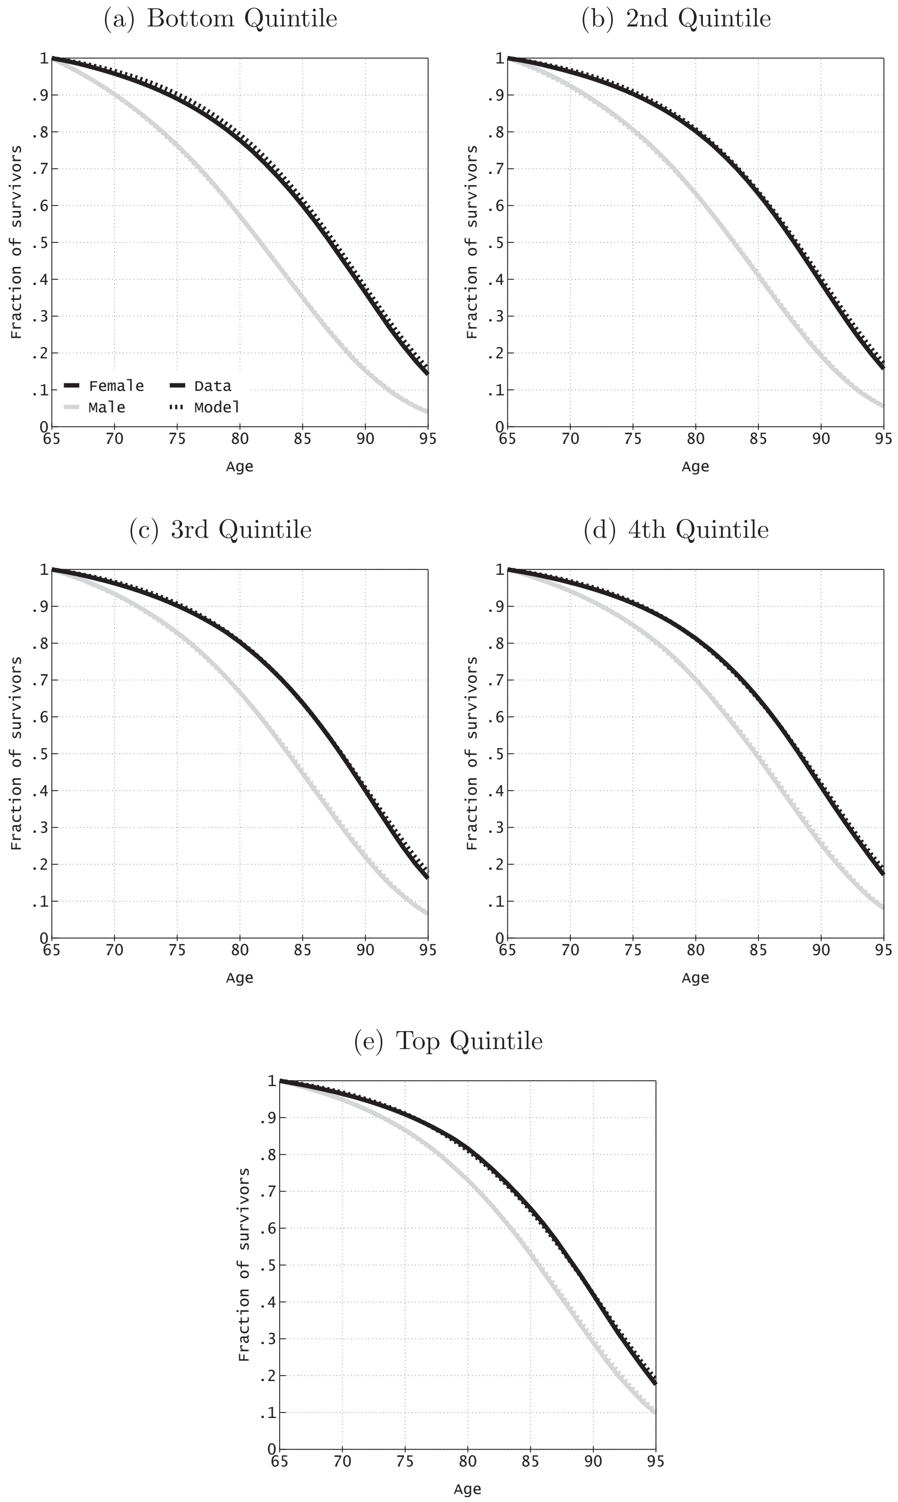

Notes: The figure compares the empirical survival curves with their simulated counterpart. The simulated curves are population-averaged measures of a life cycle simulation of 100,000 households with 1,000 bootstrapped samples.

Figure 3: Goodness of Fit of Long-term Care Use by Age

(a) Bottom Quintile

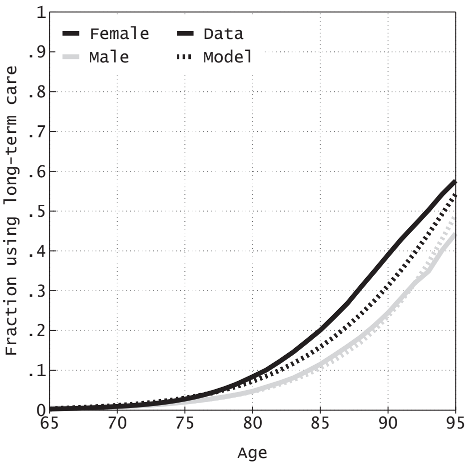

(b) 2nd Quintile

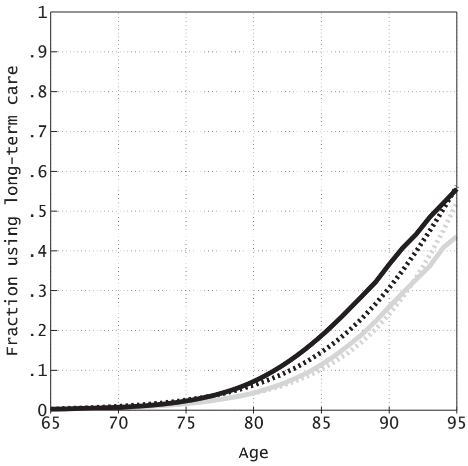

(c) 3rd Quintile

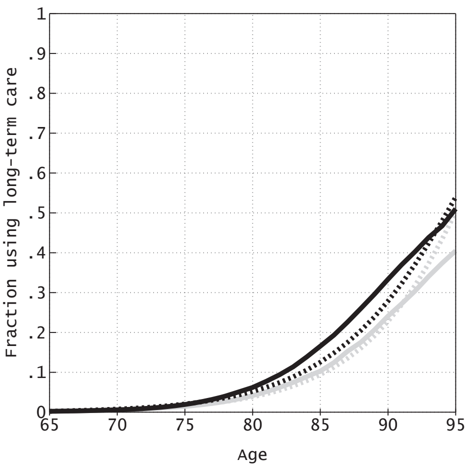

(d) 4th Quintile

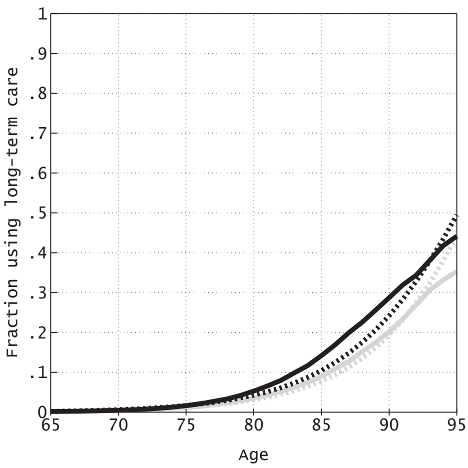

(e) Top Quintile

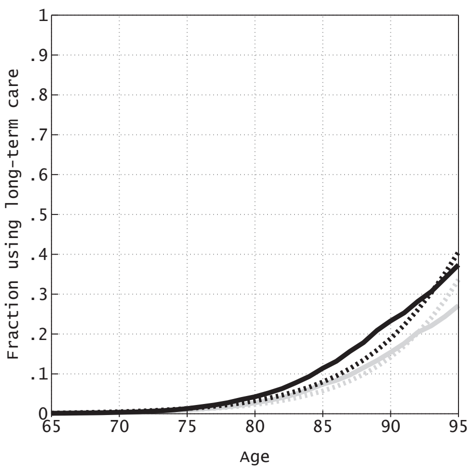

Notes: The figure compares the empirical long-term care curves with their simulated counterpart. The simulated curves are population-averaged measures of a life cycle simulation of 100,000 households with 1,000 bootstrapped samples.

B.3 Age Profile on Income

Figure 4 presents the model estimates for the age profile $\{\bar{c} + \log(\alpha_t)\}_{t=25}^{64}$. $\bar{c}$ is the fixed effect for the 1950 cohort, which we add because we want to tailor the income profile to the 1950 cohort. Figure 4 displays a familiar hump-shape (cf. Mincer, 1974): income peaks at age 55 and decreases after that. This pattern arises due to the accumulation and decumulation of human capital –working experience– over the life cycle, and households start to work less when retirement nears.

Figure 4: Estimated Age Profile on Income

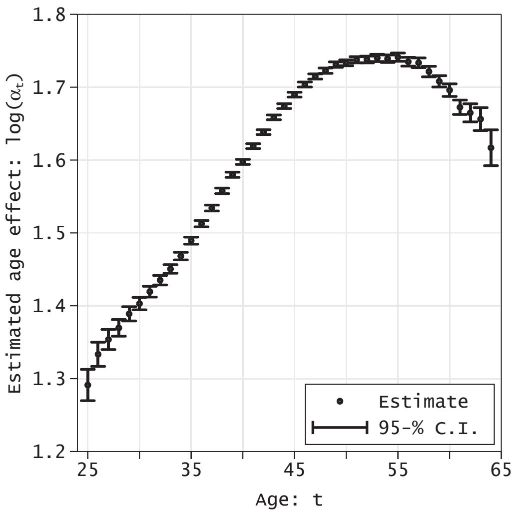

Notes: Income is measured in 0000s euros. Parameters are estimated for married households whose oldest member is younger than 65 and born after 1949. Adding $\bar{c}$ implies normalized estimates that refer to the age effect for those born in 1950. Data from the IPO 2001-2014: 77,118 households and 534,006 panel-year observations.

B.4 Income Uncertainty

We model household income dynamics as an AR(1) (canonical) process:

$$\log(y_t) = \log(\alpha_t) + \theta + \eta_t + \epsilon_t$$

$$\eta_t = \rho \cdot \eta_{t-1} + u_t$$

$$\theta \sim \mathcal{N}(0, \sigma_\theta^2); \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2); \quad u_t \sim \mathcal{N}(0, \sigma_u^2); \quad \eta_{24} = 0$$

First, we estimate the age effects $\log(\alpha_t)$ by running a fixed effects regression of log household income on age dummies, where each dummy represents a distinct effect $\log(\alpha_t)$. Next, to wash out birth cohort effects, we regress the estimate $\hat{\theta}_i$ on birth year dummies and impute the household's $\hat{\theta}_i$ to the value it would have if born in 1950. We then estimate the uncertain income component $\theta + \eta_t + \epsilon_t$ by minimum distance estimation, minimizing the squared difference between theoretical and empirical moments (cf. Storesletten et al., 2004). Because we have an auto-regressive process with a lag of one year, we match the variance and first-order auto-correlation of the income component.

The assumptions on the persistent income component imply the following process in terms of the past and current shocks:

$$\eta_t = \rho^{t-24} \cdot \eta_{24} + \sum_{j=25}^{t} \rho^{t-j} \cdot u_j + \epsilon_t, \quad t = 25, ., 64$$

from which the moments

$$\begin{aligned} \mathrm{var}(\theta + \eta_t + \epsilon_t) &= \sigma_\theta^2 + \rho^{2(t-24)} \cdot \sigma_z^2 + \sum_{j=25}^{t} \rho^{2(t-j)} \cdot \sigma_u^2 + \sigma_\epsilon^2 \\ \mathrm{cov}(\theta + \eta_t + \epsilon_t,\ \theta + \eta_{t-1} + \epsilon_{t-1}) &= \sigma_\theta^2 + \rho^{2(t-24)-1} \cdot \sigma_z^2 + \sum_{j=25}^{t} \rho^{1+2(t-j)} \cdot \sigma_u^2 \end{aligned}$$

follow, allowing us to identify the moments. Identification follows standard covariance arguments. For further details on identification, we refer to Arellano (2003).

We employ a weighted minimum distance estimator to fit these 79 moments (40 for the variances 39 for the covariances). The objective function is the sum of squared differences between the theoretical and empirical variances and co-variances. Due to the small sample considerations explained in Altonji and Segal (1996), our estimator employs the identity matrix as the weighting matrix. Hence, each moment receives the same weight in the objective function. The estimator, which minimizes the objective function, yields consistent but possibly inefficient estimates. Figure 4 and Table 2 in Section 4.1 present the estimates for the structural parameters.

Figure 5 shows the goodness-of-fit of the model estimates for the targeted moments.

Figure 5: Fit of the Income Process Before Age 65

(a) Variance of residual income (b) Auto-correlation of residual income

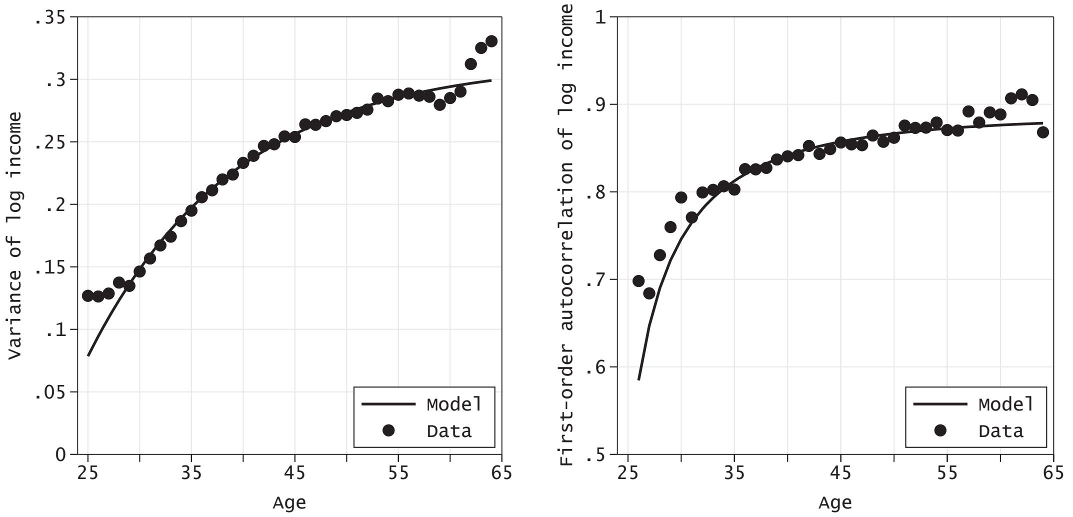

Notes: Income measured in 0000s euros. We report the parameters for married households whose oldest member is younger than age 65 and born after 1949. Data from the IPO 2001-2014: 77,118 households and 534,006 panel-year observations.

Our model matches the variance and first-order auto-correlation (closely related to first-order autocovariance) of the income shock process well. Notably, the variance of the income shock increases over time, implying more heterogeneity in income when age increases. This is important when constructing heterogeneity in asset profiles with our life-cycle model.

B.5 Replacement rates

We compute the replacement rates of survivor pensions using the IPO data restricted to households whose members are all aged 65 and over. Both members must have retirement income as their primary income source. The IPO does not distinguish between occupational pension benefits and income from privately purchased annuities (third pillar), so the replacement rate reflects both occupational and privately-arranged pension benefits. We run a fixed effects regression of log private pension income on year dummies and the family structure: being a couple, a single man, or a single woman. The exponentiated coefficient for singles gives their replacement rate. The estimates for a single man or woman are $\mathrm{rr}_m = 0.93$ (SE: 0.001) and $\mathrm{rr}_f = 0.55$ (SE: 0.005), respectively. The widow's replacement rate means that each euro of a defined pension benefit drops to 55 cents when the female spouse survives. In line with our earlier work van der Vaart et al. (2020), we report $\mathrm{rr}_m > \mathrm{rr}_f$ implied by that men were the prime earner in the households and pension benefits mostly accrued to them.

B.6 Tax Function Estimates

For general taxes, we estimate the following specification (cf. Heathcote et al., 2020):

$$\tau_G(y, \cdot) = y - \lambda \cdot y^{1-\tau},$$

which we estimate conditional upon age group (below vs. above age 65) and family structure (married vs. single).

Table 6 shows the estimates. Our estimates are in the ballpark of Heathcote et al. (2020). Using data from the Congressional Budget Office, they report $\tau \in (0.089, 0.236)$ for the U.S. between 2012-2016. λ is merely a level effect and thus does not have appropriate benchmark values.

Table 6: Parameters of the General Income Tax Function τ_G

	Couples		Singles
	Below age 65	Above age 65	Above age 65
λ	1.241	1.157	1.073
	(0.005)	(0.008)	(0.012)
τ	0.185	0.162	0.148
	(0.002)	(0.005)	(0.010)
No. households:	77,118	18,325	14,176
Panel-year observations:	534,006	101,067	64,571

Income measured in 0000s euros. Estimates for the group younger than 65 restricts to households whose oldest member is younger than 65 and born after 1949. Estimates for the group older than age 65 restricts to households whose youngest member is older than 65 and born before 1950. Standard errors (in parentheses) are clustered at the household level.

For dedicated taxes for first pillar pension (τ_{SS}) and LTC provision (τ_L), we estimate the following specification:

$$\tau_x(y,\cdot) = \alpha_{0,x} + \frac{\alpha_{1,x} - \alpha_{0,x}}{1 + e^{-\left(\frac{y-\alpha_{2,x}}{\alpha_{3,x}}\right)}}, \quad x \in \{LTC,\ SS\}, \tag{17}$$

which we estimate conditional upon age group (below vs. above age 65) and family structure (married vs. single). $\alpha_{1,x}$ represents the maximum tax amount, which is present in the Dutch system. Table 7 shows the estimation results.

Table 7: Parameters of the Dedicated Tax Functions τ_L and τ_{SS}

	Couples		Singles
	Below age 65	Above age 65	Above age 65
Pension income ($x = SS$)			
α_0	-0.255		
	(0.013)		
α_1	0.697		
	(0.002)		
α_2	3.259		
	(0.041)		
α_3	1.566		
	(0.022)		
LTC provision ($x = LTC$)			
α_0	-0.166	-0.060	-0.026
	(0.008)	(0.004)	(0.005)
α_1	0.447	0.378	0.303
	(0.001)	(0.003)	(0.003)
α_2	3.268	3.510	2.578
	(0.004)	(0.001)	(0.002)
α_3	1.599	0.872	0.618
	(0.230)	(0.021)	(0.021)
No. households:	77,118	18,325	14,176
Panel-year observations:	534,006	101,067	64,571

Income measured in 0000s euros. Estimates for the group younger than 65 are restricted to households whose oldest member is younger than 65 and born after 1949. Estimates for the group older than 65 are restricted to households whose youngest member is older than 65 and born before 1950. Standard errors are clustered at the household level (in parentheses).

B.7 Cohort Effects to the Asset Profiles

Akin to estimating the income processes before age 65, we have to deal with cohort effects to observed asset profiles. In the cross-section (a given year), older households are born in an earlier year than younger households and, due to secular income growth, have a lower labor productivity level and pension income. Because of this, asset levels of older cohorts will likely be lower. At the same time, assets of older cohorts may be higher because they include more former entrepreneurs, such as farmers. Computing age profiles of assets unconditionally upon birth cohort would consist of these undesired cohort effects.

To obtain asset profiles without cohort effects, we follow French (2005) and run specifications (2a) and (2b) with the logarithm of assets a_{it} as outcome:

$$\log(a_{it}) = \log(\alpha_{t,w}) + \theta_{i,w} + \epsilon_{it,w}, \tag{18a}$$

where i indexes a household and t is the age of the household, i.e., the age of the oldest household member. This age ranges from 65 to 100. w is a subscript to distinguish these parameters involving assets from those involving income in specifications (2a) and (2b). To wash out cohort effects, we run the following OLS regression of the predicted fixed effects on birth cohort dummies (cf. French, 2005; De Nardi et al., 2023):

$$\widehat{\theta}_{i,w} = \overline{\theta_w} + \overline{\theta}_{c,w} + \widetilde{\theta}_{i,w}, \;\; c \in \{1905, 1906, ..., 1944, 1945 - 1949\}, \tag{18b}$$

where $\overline{\theta_w}$ is the cohort effect of birth years 1945-1949, $\overline{c_w} + \overline{\theta}_{c,w}$ is the fixed effect for the other cohorts, and residual $\widetilde{\theta}_i$ is the household-specific effect excluding a cohort effect. To align with the income process before age 65 being tailored to households born in 1950, we take the cohort born between 1945 and 1949, as the reference group. In the ideal econometric scenario, we have $\overline{\theta}_{c,w} = 0$ so no cohort effects. To mimic this, we subtract the estimated cohort effect $\widehat{\overline{\theta}}_{c,w}$ from the right-hand side of (18a):

$$\log(\widehat{a}_{it}) = \log(\widehat{\alpha}_{t,w}) + \widehat{\theta}_{i,w} + \widehat{\epsilon}_{it,w} - \widehat{\overline{\theta}}_{c,w}. \tag{18c}$$

$\log(\widehat{a}_{it})$ is the predicted asset level for the household when they would be born between

1945 and 1949. To allow for distinct age patterns by marital status and lifetime income, we run the regressions for these groups separately.

We exponentiate the assets to get the asset level that is cleaned from cohort effects. While the regression omits zero assets, we re-include them in the 'cleaned' profiles; negative assets and assets above €2,500,000 are dropped.[16][17]

Figure 6 shows median asset profiles before and after we control for birth cohort effects. Each separate line represents a different birth cohort, depending on the age in 2006. The left panels a. and c., i.e., the raw data, reveal that birth cohort effects are strong, particularly for married households with high lifetime income. Those households have more assets if they are born earlier. Furthermore, within birth cohorts, there seems to be a strong time trend, induced by the period of financial crisis that is part of our observational window.

Using (18c), a birth cohort effect is controlled for in panels b. and d.. This reverses the differences between cohorts: the youngest cohorts hold most assets and asset profiles of different cohorts nicely overlap. Also, year trends are less pronounced. As a consequence, we observe households decumulating asset holdings over time. The asset profiles in Figure 1 in Section 5, which we target, are the data from panels b. and d. unconditional upon birth cohort.

C Second-stage Estimates

C.1 Standard Errors of Estimated Preference Parameters

We compute standard errors of $\widehat{\boldsymbol{\delta}}$ by using a matrix $\boldsymbol{D}$ that measures the responsiveness of each moment condition to slightly changing the parameter estimate. Specifically, $\boldsymbol{D}$ is a $k \times 3$ dimensional matrix where the k-th row contains the derivative of the k-th moment condition: $\frac{\partial \left(M_k^d - M_k^s(\hat{\chi}, \boldsymbol{\delta})\right)}{\partial \boldsymbol{\delta}}$. The variance-covariance $\boldsymbol{V}$ of estimator $\widehat{\boldsymbol{\delta}}$ is documented in De Nardi et al. (2010): $\boldsymbol{V} = \left(\boldsymbol{D}'\boldsymbol{D}\right)^{-1}\left(\boldsymbol{D}'\boldsymbol{S}\boldsymbol{D}\right)\left(\boldsymbol{D}'\boldsymbol{D}\right)^{-1}$, where $\boldsymbol{S}$ is the empirical

[16] We drop 0.9% of the households and 2.6% of the panel-year observations because of these restrictions.

[17] We also tried Deaton-Paxson dummies, but identifying the effects suffers heavily from multicollinearity. Also, taking levels as outcome could not properly control for many zero assets in the data.

Figure 6: Asset Profiles Before and After Controlling for Birth Cohort Fixed Effects

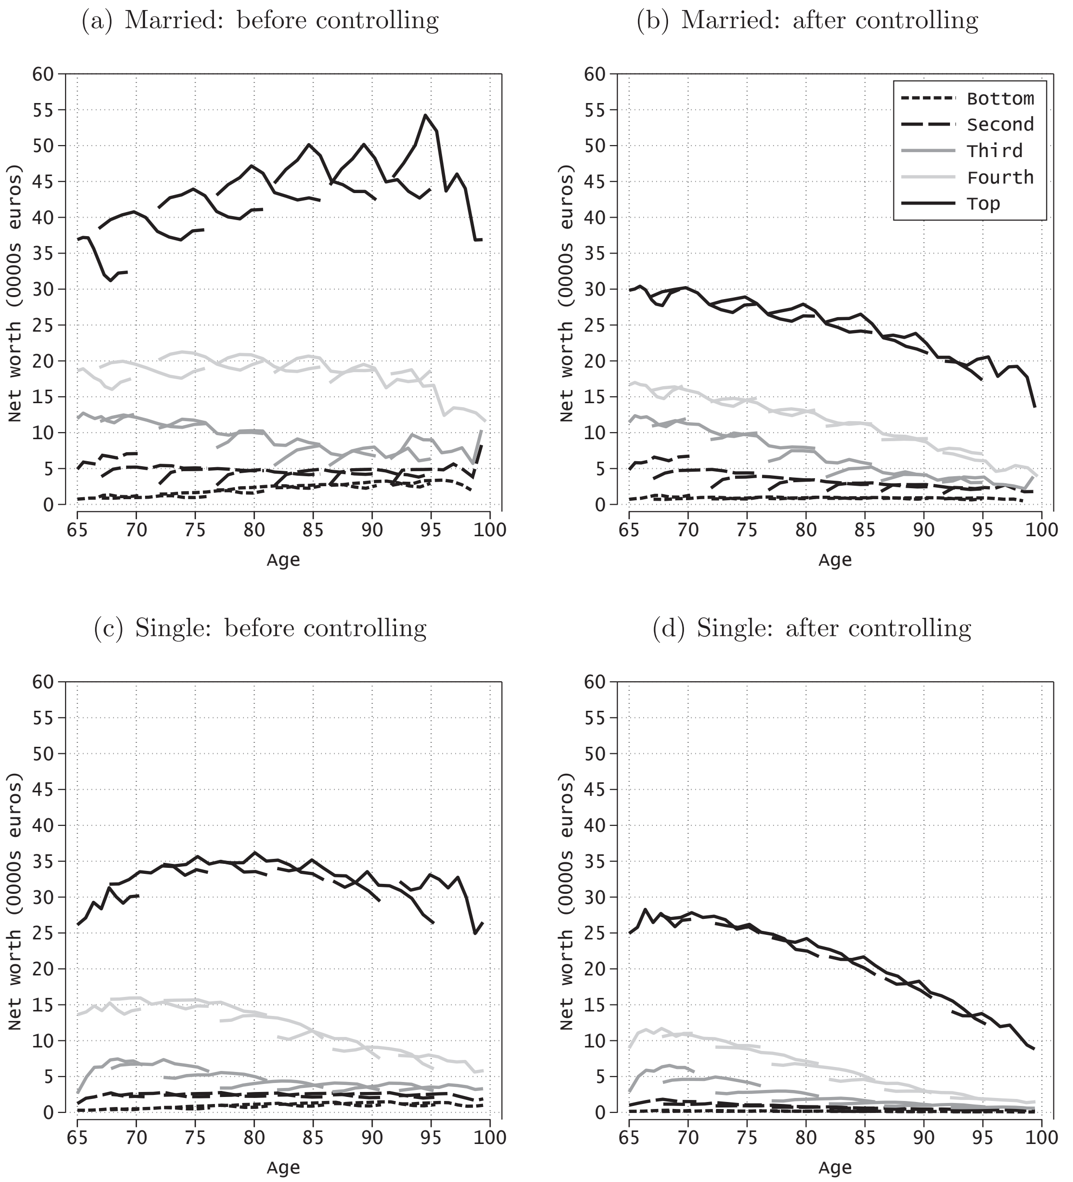

Notes: Each line represents the asset profile conditional upon birth cohort and income quintile. We distinguish seven birth cohorts based on the age of the household in 2006: younger than 65; aged 65-69; aged 70-74; aged 75-79; aged 80-84; aged 85-89; and aged 90 and over.

variance-covariance matrix regarding the data moments. We compute $\boldsymbol{D}$ numerically.